

Ground state connectivity of local Hamiltonians

Sevag Gharibian*

Jamie Sikora†

February 20, 2015

Abstract

The study of ground state energies of local Hamiltonians has played a fundamental role in quantum complexity theory. In this paper, we take a new direction by introducing the physically motivated notion of “ground state connectivity” of local Hamiltonians, which captures problems in areas ranging from quantum stabilizer codes to quantum memories. We show that determining how “connected” the ground space of a local Hamiltonian is can range from QCMA-complete to NEXP-complete. As a result, we obtain a natural QCMA-complete problem, a goal which has generally proven difficult since the conception of QCMA over a decade ago. Our proofs rely on a new technical tool, the Traversal Lemma, which analyzes the Hilbert space a local unitary evolution must traverse under certain conditions. We show that this lemma is tight up to a polynomial factor with respect to the length of the unitary evolution in question.

1 Introduction

Over the last fifteen years, the merging of condensed matter physics and computational complexity theory has given rise to a new field of study known as *quantum Hamiltonian complexity* [Osb12, GHLS14]. The cornerstone of this field is arguably Kitaev’s [KSV02] quantum version of the Cook-Levin theorem [Coo72, Lev73], which says that the problem of estimating the ground state energy of a local Hamiltonian is complete for the class Quantum Merlin Arthur (QMA), where QMA is a natural generalization of NP. Here, a k -local Hamiltonian is an operator $H = \sum_i H_i$ acting on n qubits, such that each local Hermitian constraint H_i acts non-trivially on k qubits. The *ground state energy* of H is simply the smallest eigenvalue of H , and the corresponding eigenspace is known as the *ground space* of H .

Kitaev’s result spurred a long line of subsequent works on variants of the ground energy estimation problem (see, e.g. [Osb12, GHLS14] for surveys), known as the k -local Hamiltonian problem (k -LH). For example, Oliveira and Terhal showed that LH remains QMA-complete in the physically motivated case of qubits arranged on a 2D lattice [OT08]. Bravyi and Vyalı proved [BV05] that the *commuting* variant of 2-LH is in NP. More recently, the complexity of the version of 2-LH in which large positive and negative weights on local terms are allowed was characterized by Cubitt and Montanaro [CM13] in a manner analogous to Schaeffer’s dichotomy theorem for Boolean satisfiability [Sch78]. Thus, k -LH has served as an excellent “benchmark” problem for delving into the complexity of problems encountered in the study of local Hamiltonians. Yet, one can also ask about the *properties of the ground space* itself. For example, is it topologically ordered? Can we evaluate local observables against it [Osb12]? It is this direction which we pursue in this paper.

*Simons Institute for the Theory of Computing and Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA 94720, USA, and Department of Computer Science, Virginia Commonwealth University, Richmond, VA 23284, USA. Email: sgharibian@vcu.edu.

†Centre for Quantum Technologies, National University of Singapore, Singapore 117543. Email: cqtjwjs@nus.edu.sg.

Specifically, in this paper we define a notion of *connectivity* of the ground space of H , which roughly asks: Given ground states $|\psi\rangle$ and $|\phi\rangle$ of H as input, are they “connected” through the ground space of H ? Somewhat more formally, we have (see Section 2 for a formal definition):

Definition 1.1 (Ground State Connectivity (GSCON) (informal)). *Given as input a local Hamiltonian H and ground states $|\psi\rangle$ and $|\phi\rangle$ of H (specified via quantum circuits), as well as parameters m and l , does there exist a sequence of l -qubit unitaries $(U_i)_{i=1}^m$ such that:*

1. $(|\psi\rangle \text{ mapped to } |\phi\rangle) U_m \cdots U_1 |\psi\rangle \approx |\phi\rangle$, and
2. $(\text{intermediate states in ground space}) \forall i \in [m], U_i \cdots U_1 |\psi\rangle \text{ is in the ground space of } H?$

In other words, GSCON asks whether there exists a sequence of m unitaries, each acting on (at most) l qubits, mapping the initial state $|\psi\rangle$ to the final state $|\phi\rangle$ *through* the ground space of H . We stress that the parameters m (i.e. number of unitaries) and l (i.e. the locality of each unitary) are key; as we discuss shortly, depending on their setting, the complexity of GSCON can vary greatly.

Physics Motivation. The original inspiration for this work came from a recently active area in classical complexity theory on *reconfiguration* problems (see *Previous work* below for details). For example, the reconfiguration problem for 3SAT asks: Given a 3SAT formula ϕ and satisfying assignments x and y for ϕ , does there exist a sequence of bit flips mapping x to y , such that each intermediate assignment encountered is also a satisfying assignment for ϕ ? Although the classical study of reconfiguration problems is arguably mostly interesting from a theoretical perspective, its quantum variant (i.e. GSCON) turns out to be physically relevant. To illustrate, we now discuss connections to *quantum memories* and *stabilizer codes*.

Quantum memories. A key challenge in building quantum computers is the implementation of long-lived qubit systems. In low-temperature systems, one approach is to encode a qubit in the ground state of a gapped Hamiltonian with a degenerate ground space. Here, the degeneracy ensures the qubit has at least two basis states, logical $|\tilde{0}\rangle$ and $|\tilde{1}\rangle$, and the gap ensures that external noise does not (easily) take a ground state out of the ground space. However, this is not sufficient — although environmental noise may not take the state *out* of the ground space, it can still alter the state *within* the ground space (e.g. inadvertently map $|\tilde{0}\rangle$ to $|\tilde{1}\rangle$). Thus, making the typical assumption that errors act locally, it should ideally not be possible for $|\tilde{0}\rangle$ to be mapped to $|\tilde{1}\rangle$ through the ground space via a sequence of local operations. This is precisely the principle behind Kitaev’s toy chain model [Kit01], and the motivation behind the toric code [Kit03] (see also [KL09]). This notion of how “robust” a quantum memory is can thus be phrased as an instance of GSCON: Given a gapped Hamiltonian H , a ground state $|\psi\rangle$ to which the quantum memory is initialized, and an undesired ground state $|\phi\rangle$, is there a sequence of local errors mapping the state of our quantum memory through the ground space from $|\psi\rangle$ to $|\phi\rangle$?

Stabilizer codes. Roughly, a stabilizer code [Got97] is a quantum error-correcting code defined by a set of commuting Hermitian operators, $S = \{G_1, \dots, G_k\}$, such that $G_i \neq -I$ and $\|G_i\|_\infty \leq 1$ for all $G_i \in S$. The *codespace* for S is the set of all $|\psi\rangle$ satisfying $G_i |\psi\rangle = |\psi\rangle$ for all $i \in [k]$. In other words, defining G_i^+ as the projection onto the $+1$ eigenspace of G_i , the codespace is the ground space of the positive semidefinite Hamiltonian $H := \sum_{i=1}^k (I - G_i^+)$. Typically, errors are assumed to occur on a small number of qubits at a time; with this assumption in place, the following is a special case of GSCON: Given H and codewords $|\psi\rangle$ and $|\phi\rangle$, does there exist a sequence of at

most m local errors mapping $|\psi\rangle$ to $|\phi\rangle$, such that the entire error process is undetectable, i.e. each intermediate state remains in the codespace?

Results. Having motivated GSCON, we now informally state our results.

Theorem 1.2 (See Theorem 5.1 for a formal statement). *GSCON for polynomially large m (i.e. for polynomially many local unitaries U) and $l = 2$ (i.e. 2-qubit unitaries) is QCMA-complete.*

Here, QCMA is QMA except with a classical prover [AN02]. See Section 2 for a formal definition. Theorem 1.2 says that determining whether there exists a polynomial-size quantum circuit mapping $|\psi\rangle$ to $|\phi\rangle$ through the ground space of H is QCMA-complete.

Theorem 1.3 (See Theorem 6.1 for a formal statement). *GSCON for exponentially large m (i.e. for exponentially many local unitaries U) and $l = 1$ (i.e. 1-qubit unitaries) is PSPACE-complete.*

Theorem 1.3 says that determining whether there exists an exponential length sequence of 1-qubit unitaries mapping $|\psi\rangle$ to $|\phi\rangle$ through the ground space of H is PSPACE-complete.

Finally, in Section 7 we define a succinct variant of GSCON, called SUCCINCT GSCON, in which the Hamiltonian H has a succinct circuit description, and the initial and final states $|\psi\rangle$ and $|\phi\rangle$ are product states. We show:

Theorem 1.4 (See Theorem 7.4 for a formal statement). *SUCCINCT GSCON for exponentially large m (i.e. for exponentially many local unitaries U) and $l = 1$ (i.e. 1-qubit unitaries) is NEXP-complete.*

As Theorem 1.4 follows from techniques similar to Theorems 1.2 and 1.3, we give only a proof sketch of it in Section 7.

We remark that the choices of m and l above are key to our results. For example, Theorem 1.2 holds for any constant $l \geq 2$ (see remarks after its proof); however, for $l \in \omega(\log N)$ (for N the input size) the problem is likely no longer in QCMA, as the prover cannot send a classical description of each local unitary. Similarly, attempting to extend Theorem 1.3 by setting $l = 2$ appears problematic, as then any intermediate state in the unitary evolution seems to require exponential space to represent. This latter problem is, however, in NEXP. We thus conjecture that it is actually NEXP-complete.

Proof techniques. Our results rely on a new technical lemma called the Traversal Lemma, as well as the use of ϵ -nets and ϵ -pseudo-nets (also known as *improper covering sets*). We now outline the proof techniques behind Theorem 5.1 (QCMA-completeness) in more detail; using similar ideas, Theorems 6.1 (PSPACE-completeness) and 7.4 (NEXP-completeness) follow analogously.

Specifically, we outline both QCMA-hardness and containment in QCMA. Beginning with the former, the central idea behind the construction is as follows. Let V be an arbitrary QCMA verification circuit, and let H' be the local Hamiltonian obtained from V via Kitaev's circuit-to-Hamiltonian construction [KSV02] (see Lemma 2.4 for Kempe and Regev's 3-local version [KR03]). Then, we design the input Hamiltonian H to GSCON so that "traversing its ground space" is equivalent to simulating the following protocol: Starting from the all-zeroes state, prepare the ground state of H' (which can be done efficiently since V is a QCMA circuit), and subsequently flip a set of special qubits called *GO* qubits. This latter step "activates" the check Hamiltonian H , which now "verifies" that the ground state prepared is indeed correct. Finally, uncompute the ground state to arrive at a target state of all-zeroes except in the *GO* register, which is now set to all ones.

To prove correctness of this construction, our main technical tool is a new lemma we call the *Traversal Lemma* (Lemma 4.2), which analyzes the Hilbert space a local unitary evolution must traverse in certain settings. Specifically, define two states $|\psi\rangle$ and $|\phi\rangle$ as *k-orthogonal* if for any k -local unitary U , we have $\langle\phi|U|\psi\rangle = 0$. In other words, any application of a k -local unitary leaves $|\psi\rangle$ and $|\phi\rangle$ orthogonal. Then, the Traversal Lemma says that for k -orthogonal states $|\psi\rangle$ and $|\phi\rangle$, if we wish to map $|\psi\rangle$ to $|\phi\rangle$ via a sequence of k -local unitaries, then at some step in this evolution we must leave the space spanned by $|\psi\rangle$ and $|\phi\rangle$, i.e. we must have “large” overlap with the orthogonal complement of $|\psi\rangle$ and $|\phi\rangle$. To prove the Traversal Lemma, we use a combination of the Gentle Measurement Lemma of Winter [Win99] and an idea inspired by the quantum Zeno effect.

As the Traversal Lemma is a key technical contribution of this paper, we also study its properties further (i.e. independently of its application to our complexity theoretic results). For example, we show the lemma is tight up to a polynomial factor in the number of unitaries. To do so, we give a pair of 3-orthogonal states $|\psi\rangle, |\phi\rangle$ with the following property: For any $0 < \Delta < 1/2$, we construct a carefully selected sequence of $O(1/\Delta^2)$ 2-local unitaries mapping $|\psi\rangle$ to $|\phi\rangle$, such that at any point in this mapping, the overlap with the orthogonal complement of $|\psi\rangle$ and $|\phi\rangle$ is at most Δ . We also delve further into the study of k -orthogonality, including giving an intuitive characterization of the notion.

Finally, containment of GSCON in QCMA is shown via a simple and natural verification procedure, wherein the prover sends a classical description of the local unitaries $\{U_i\}$, and the verifier prepares many copies of the starting, final, and all intermediate states and checks that all required properties hold. To make this rigorous, we construct an ϵ -pseudo-net, which allows us to easily discretize the space of d -dimensional unitary operators for any $d \geq 2$. Such pseudo-nets come with a tradeoff: On the negative side, they contain non-unitary operators. On the positive side, they are not only straightforward to construct, but more importantly, they have the following property: Given any element A in the pseudo-net, there are efficient *explicit* protocols for checking if A is close to unitary, and if so, for “rounding” it to such a unitary.

Previous work. To the best of our knowledge, our work is the first to study reconfiguration in the quantum setting. In contrast, in the classical setting, such problems have recently received much attention. In particular, our work was inspired by the paper of Gopalan, Kolaitis, Maneva, and Papadimitriou [GKMP06], which shows that determining whether two solutions x and y of a Boolean formula are connected through the solution space is either in P or is PSPACE-complete, depending on the constraint types allowed in the formula. (Note: A minor error in Reference [GKMP06] was recently corrected in the work of Schwerdtfeger [Sch13].) More recently, Mouawad, Nishimura, Pathak and Raman [MNPR14] studied the variant of this problem in which one seeks the *shortest* possible Boolean reconfiguration path; they show this problem is either in P, NP-complete, or PSPACE-complete. In this sense, our definition of GSCON can be thought of as a quantum generalization of the problem studied in Reference [MNPR14]. More generally, since the work of Reference [GKMP06], a flurry of papers have appeared studying reconfiguration for problems ranging from Boolean satisfiability to vertex cover to graph coloring [CvdHJ08, BC09, BJL⁺11, CvdHJ11, FHHH11, IDH⁺11, Bon12, IKD12, IKOZ12, KMM12, Sch13, BB13, MNR⁺13, MNPR14, MNR14].

Significance to complexity theory. We now discuss the motivation behind GSCON from a complexity theoretic perspective. We begin by focusing on QCMA, which is a natural class satisfying $\text{MA} \subseteq \text{QCMA} \subseteq \text{QMA}$. Although QCMA was introduced over a decade ago by Aharonov and Naveh [AN02], we still have an unfortunately small number of complete problems for it. In particular, to the best of our knowledge, the following is an exhaustive list:

- Does a given local Hamiltonian have an efficiently preparable ground state [WJB03]?
- Does a given quantum circuit act almost as the identity on computational basis states [WJB03]?
- Given a braid, can it be conjugated by another braid from a given class such that the Jones polynomial of its plat closure is nearly maximal [WY08]?
- Given a continuous-time classical random walk on a restricted class of graphs, and time T , do there exist vertices i and j such that the difference of the probabilities of being at i and j is at least $c \cdot \exp(-\mu T)$ [JW06]?
- Given a quantum circuit C accepting a non-empty monotone set, what is the smallest Hamming weight string accepted by C [GK12]?

In this regard, the pursuit of natural complete problems for QCMA has arguably proven rather difficult. Our results add a new, physically-motivated problem to the short list of QCMA-complete problems.

Second, a common focus in quantum complexity theory has been the problem of estimating the ground state energy of a given local Hamiltonian (see, e.g. [GLSW14] for a survey). However, less attention has been given to the complexity of determining other properties of local Hamiltonians. For example, Brown, Flammia, and Schuch showed [BFS11] that computing the ground state degeneracy and density of states for a local Hamiltonian is $\#\text{BQP}$ -complete. Gharibian and Kempe showed [GK12] that determining the smallest subset of interaction terms of a given local Hamiltonian which yields a high energy ground space is $\text{cQ-}\Sigma_2$ -complete. Ambainis has shown [Amb14] (among other results) that evaluating local observables against a local Hamiltonian is $\text{PQMA}^{\log n}$ -complete, and that determining the spectral gap of a local Hamiltonian is in $\text{PQMA}^{\log n}$. Continuing in this vein, our work initiates a new direction of study regarding properties of local Hamiltonians beyond estimating the ground state energy, namely the study of ground state connectivity.

Finally, regarding the use of our proof techniques in the study of quantum algorithms and verification procedures, we hope the Traversal Lemma may prove useful in its own right. For example, in quantum adiabatic algorithms, it is often notoriously difficult to understand how a quantum state evolves in time from an easy-to-prepare initial state to some desired final state. The Traversal Lemma gives us a tool for studying the behaviour of such evolutions, playing a crucial role in our analysis here. We remark, however, that in quantum adiabatic evolution, the Hamiltonian itself changes with time, whereas here our Hamiltonian is fixed and we apply local unitary gates to our quantum state.

Organization. This paper is organized as follows. In Section 2, we state relevant notation, definitions, and useful known results. Section 3 constructs ϵ -nets and ϵ -pseudo-nets over unitary operators, which are used in Sections 5, 6 and 7 for showing containment of GSCON in QCMA, PSPACE, and NEXP, respectively. Section 4 introduces the notion of k -orthogonality and states and proves the Traversal Lemma, which is used in Sections 5, 6, and 7 to show QCMA-hardness, PSPACE-hardness, and NEXP-hardness of GSCON. Section 4.1.1 shows our result regarding tightness of the Traversal Lemma and Section 4.1.2 studies the properties of k -orthogonality further. We conclude and state open problems in Section 8.

2 Preliminaries

Notation. The notation $:=$ is used to indicate a definition. Given $x \in \{0, 1\}^n$, $|x\rangle \in (\mathbb{C}^2)^{\otimes n}$ denotes the computational basis state labeled by x . For a vector $|v\rangle$, define its Euclidean norm

as $\| |v\rangle \|_2 := (\sum_i |v_i|^2)^{1/2}$ and its infinity norm as $\| |v\rangle \|_\infty := \max_i |v_i|$. For complex Euclidean space \mathcal{X} , let $\mathcal{L}(\mathcal{X})$, $\text{Herm}(\mathcal{X})$ and $\mathcal{U}(\mathcal{X})$ denote the sets of linear, Hermitian and unitary operators acting on \mathcal{X} , respectively. We use the following matrix norms: $\|A\|_{\max} := \max_{ij} |A(i,j)|$, the spectral norm $\|A\|_\infty := \max\{\|A|v\rangle\|_2 : \| |v\rangle \|_2 = 1\}$, and trace norm $\|A\|_{\text{tr}} := \text{Tr} \sqrt{A^\dagger A}$. The Hilbert-Schmidt or trace inner product between operators A and B is $\langle A, B \rangle := \text{Tr}(A^\dagger B)$. The set of natural numbers is \mathbb{N} , and $[m] := \{1, \dots, m\}$. Throughout this paper, we treat the local dimension d of quantum systems as a constant.

Definitions. We now formally define the problem studied in this paper. (To ease parsing of the definition, the input parameters are highlighted in maroon online.)

Definition 2.1 (Ground State Connectivity (GSCON($H, k, \eta_1, \eta_2, \eta_3, \eta_4, \Delta, l, m, U_\psi, U_\phi$))).

Input parameters:

1. k -local Hamiltonian $H = \sum_i H_i$ acting on n qubits with $H_i \in \text{Herm}((\mathbb{C}^2)^{\otimes k})$ satisfying $\|H_i\|_\infty \leq 1$.
2. $\eta_1, \eta_2, \eta_3, \eta_4, \Delta \in \mathbb{R}$, and integer $m \geq 0$, such that $\eta_2 - \eta_1 \geq \Delta$ and $\eta_4 - \eta_3 \geq \Delta$.
3. Polynomial size quantum circuits U_ψ and U_ϕ generating “starting” and “target” states $|\psi\rangle$ and $|\phi\rangle$ (starting from $|0\rangle^{\otimes n}$), respectively, satisfying $\langle \psi | H | \psi \rangle \leq \eta_1$ and $\langle \phi | H | \phi \rangle \leq \eta_1$.

Output:

1. If there exists a sequence of l -local unitaries $(U_i)_{i=1}^m \in \mathcal{U}(\mathbb{C}^2)^{\times m}$ such that:
 - (a) (Intermediate states remain in low energy space) For all $i \in [m]$ and intermediate states $|\psi_i\rangle := U_i \cdots U_2 U_1 |\psi\rangle$, one has $\langle \psi_i | H | \psi_i \rangle \leq \eta_1$, and
 - (b) (Final state close to target state) $\|U_m \cdots U_1 |\psi\rangle - |\phi\rangle\|_2 \leq \eta_3$,
then output YES.
2. If for all l -local sequences of unitaries $(U_i)_{i=1}^m \in \mathcal{U}(\mathbb{C}^2)^{\times m}$, either:
 - (a) (Intermediate state obtains high energy) There exists $i \in [m]$ and an intermediate state $|\psi_i\rangle := U_i \cdots U_2 U_1 |\psi\rangle$, such that $\langle \psi_i | H | \psi_i \rangle \geq \eta_2$, or
 - (b) (Final state far from target state) $\|U_m \cdots U_1 |\psi\rangle - |\phi\rangle\|_2 \geq \eta_4$,
then output NO.

A few remarks are in order. First, in the Hamiltonian complexity literature the gap size Δ for energy levels of local Hamiltonians is often taken to be inverse polynomial. Some of our results require this gap to be exponentially small. Allowing Δ to be specified as input thus allows us to precisely formulate such results. Second, the circuits U_ψ and U_ϕ are assumed to be given in terms of 1 and 2-qubit unitary gates. Third, all input parameters are specified with rational entries, each using $O(\text{poly}(n))$ bits of precision.

For completeness, we next give a formal definition of the complexity class QCMA (also known as Merlin-Quantum-Arthur (MQA) [Wat09]).

Definition 2.2 (QCMA). A promise problem $A = (A_{\text{yes}}, A_{\text{no}})$ is in QCMA if and only if there exist polynomials p, q and a polynomial-time uniform family of quantum circuits $\{Q_n\}$, where Q_n takes as input a string $x \in \Sigma^*$ with $|x| = n$, a classical proof $y \in \{0, 1\}^{\otimes p(n)}$, and $q(n)$ ancilla qubits in state $|0\rangle^{\otimes q(n)}$, such that:

- (Completeness) If $x \in A_{\text{yes}}$, then there exists a proof $y \in \{0, 1\}^{\otimes p(n)}$ such that Q_n accepts (x, y) with probability at least $2/3$.
- (Soundness) If $x \in A_{\text{no}}$, then for all proofs $y \in \{0, 1\}^{\otimes p(n)}$, Q_n accepts (x, y) with probability at most $1/3$.

Useful known results. We next state known results which prove useful in this paper. The first of these is the Gentle Measurement Lemma of Winter [Win99]; the specific variant we state below is Lemma 9.4.2 from the textbook of Wilde [Wil13].

Lemma 2.3 (Gentle Measurement Lemma [Win99], as stated in Lemma 9.4.2 of [Wil13]). *Let $\rho \in \mathcal{L}(\mathbb{C}^d)$ be a density operator and $O \preceq \Lambda \preceq I$ a measurement operator for $\Lambda \in \mathcal{L}(\mathbb{C}^d)$, such that $\text{Tr}(\Lambda\rho) \geq 1 - \epsilon$. Then, $\left\| \rho - \sqrt{\Lambda}\rho\sqrt{\Lambda} \right\|_{\text{tr}} \leq 2\sqrt{\epsilon}$.*

We next recall Kempe and Regev's 3-local circuit-to-Hamiltonian construction [KR03], which maps a given quantum circuit $V = V_L \cdots V_1$ (where each V_i is at most 2-local) acting on a *proof* register (register A) and *ancilla* register (register B) to a 3-local Hamiltonian H acting on $A \otimes B \otimes C$, where C is a *clock* register (represented in unary). The precise details of the construction are not necessary for this work; rather, we require only the following key property of H . Define the *history state* for arbitrary proof $|\psi\rangle$ in register A as

$$|\psi_{\text{hist}}\rangle := \frac{1}{\sqrt{L+1}} \sum_{i=0}^L V_i \cdots V_1 |\psi\rangle_A \otimes |0\rangle_B \otimes |i\rangle_C. \quad (1)$$

Then, the question of whether V accepts $|\psi\rangle$ is related to the smallest eigenvalue of H as follows.

Lemma 2.4 (Kempe and Regev [KR03]). *Kempe and Regev's construction maps a quantum circuit V to a 3-local Hamiltonian H with parameters α and β satisfying:*

- If there exists a proof $|\psi\rangle$ accepted by V with probability at least $1 - \epsilon$, then $|\psi_{\text{hist}}\rangle$ achieves

$$\text{Tr}(H |\psi_{\text{hist}}\rangle \langle \psi_{\text{hist}}|) \leq \alpha := \epsilon/(L+1).$$

- If V rejects all proofs $|\psi\rangle$ with probability at least $1 - \epsilon$, then the smallest eigenvalue of H is at least $\beta \in \Omega\left(\frac{1}{L^3}\right)$.

We next discuss the classical reconfiguration problem for Boolean formulae known as (s, t) -Connectivity (denoted s, t -CONN, for short).

Definition 2.5 ((s, t) -CONN). *Given a Boolean 3-CNF formula ϕ and solutions $x, y \in \{0, 1\}^n$ to ϕ , does there exist a sequence of strings $(x_i)_{i=1}^m$ such that*

1. $x_1 = x$ and $x_m = y$, and
2. for all $i \in [m]$, the Hamming distance between x_i and x_{i+1} is at most 1, and
3. for all $i \in [m]$, x_i is a solution to ϕ ?

Theorem 2.6 ([GKMP06]). *s, t -CONN is PSPACE-complete.*

Finally, we state a few useful norm inequalities. For arbitrary complex unit vectors $|v\rangle$ and $|w\rangle$ (see, e.g., Equation 1.33 of Reference [Gha13]):

$$\| |v\rangle\langle v| - |w\rangle\langle w| \|_{\text{tr}} = 2\sqrt{1 - |\langle v|w\rangle|^2} \leq 2\| |v\rangle - |w\rangle \|_2. \quad (2)$$

For arbitrary (not necessarily normalized) complex vectors, we have:

$$\| |v\rangle\langle v| - |w\rangle\langle w| \|_{\text{F}} \leq (\| |v\rangle \|_2 + \| |w\rangle \|_2) \| |v\rangle - |w\rangle \|_2. \quad (3)$$

Proof. We use the triangle inequality and the fact that $\| |a\rangle\langle b| \|_{\text{F}} = \| |a\rangle \|_2 \| |b\rangle \|_2$ (seen by expanding the definition of $\| |a\rangle\langle b| \|_{\text{F}}$) to obtain:

$$\begin{aligned} \| |v\rangle\langle v| - |w\rangle\langle w| \|_{\text{F}} &\leq \| |v\rangle\langle v| - |v\rangle\langle w| \|_{\text{F}} + \| |v\rangle\langle w| - |w\rangle\langle w| \|_{\text{F}} \\ &= \| |v\rangle \|_2 (\| |v\rangle - |w\rangle \|_2) + \| (|v\rangle - |w\rangle) \langle w| \|_{\text{F}} \\ &= (\| |v\rangle \|_2 + \| |w\rangle \|_2) \| |v\rangle - |w\rangle \|_2. \end{aligned}$$

□

3 Nets and pseudo-nets over unitary operators

In order to show containment of GSCON in the complexity classes of interest, we require nets with respect to spectral norm over unitary operators. In this section, we give two types of nets: (1) An ϵ -net over single qubit unitaries (Lemma 3.1), and (2) an ϵ -pseudo-net over unitaries of any dimension $d \geq 2$ (Lemma 3.3). The former is used in Lemma 6.3 (containment in PSPACE) and Lemma 7.6 (containment in NEXP), and consists strictly of unitary operators. The latter is used in Lemma 5.3 (containment in QCMA), and is a relaxation of a net in that it contains *non-unitary* operators; this relaxed definition, however, allows for a straightforward construction in dimensions greater than two. Note that having an exact net helps make the analysis in the proof of Lemma 6.3 easier, explaining why we use both kinds of nets. We begin with a simple single-qubit ϵ -net construction.

Lemma 3.1. *For any $0 < \epsilon \leq 1$, there exists an ϵ -net with respect to the spectral norm over $\text{U}(\mathbb{C}^2)$ of size $O(\epsilon^{-8})$. Moreover, given the index i of any element U_i in the net, U_i can be computed in time $O(\log^2(1/\epsilon))$.*

The proof is given in Appendix A, and relies on a simple characterization of single qubit unitaries. For larger dimensions $d > 2$, however, we are unaware of a similar characterization. Thus, for $d > 2$ we construct¹ an ϵ -pseudo-net. Intuitively, a pseudo-net over unitary operators contains matrices which are close to, but not necessarily, unitary. However, to aid in its use, it has two important properties: First, we give an efficient “check” procedure C such that, for any unitary U , there exists a net element M satisfying $\|U - M\|_{\infty} \leq \epsilon$ and such that M is accepted by C . Second, we give an efficient “rounding” procedure R such that if net element M is accepted by C , then R rounds M to a unitary U satisfying $\|U - M\|_{\infty} \leq \epsilon$.

Definition 3.2 (ϵ -pseudo-net). *Let $S \subseteq \text{L}(\mathbb{C}^d)$. Then, we call $S' \subseteq \text{L}(\mathbb{C}^d)$ an ϵ -pseudo-net over S if there exist $O(\text{poly}(d))$ -time algorithms C (for checking) and R (for rounding) taking as input $M \in \text{L}(\mathbb{C}^d)$ such that:*

¹It was pointed out to us by an anonymous referee that there is an alternative way to construct an ϵ -net over unitary operators with $d > 2$, which can be used in place of our pseudo-net here. Namely, one casts a net over the set of Hermitian operators H satisfying $\|H\|_{\infty} \leq \pi$, and subsequently exponentiates the items in the net.

1. (Checking) $\forall M \in S$, there exists $M' \in S'$ such that C accepts M' and $\|M - M'\|_\infty \leq \epsilon$.
2. (Rounding) $\forall M' \in S'$, if C accepts M' , then algorithm R maps M' to $M \in S$ such that $\|M - M'\|_\infty \leq \epsilon$.

We now show that there is a straightforward way to construct an ϵ -pseudo-net over $S = \mathcal{U}(\mathbb{C}^d)$ for any $d \geq 2$. The ideas here are based on a standard construction for nets over unitary operators, as used in Reference [PGA⁺11] and detailed further in Lemma 7.13 of Reference [Gha13]; this standard construction is, however, inherently non-explicit. Thus, we adapt it here as necessary to obtain an *explicit* ϵ -pseudo-net.

Lemma 3.3. *For any $0 < \epsilon < 1$, there exists a set $N \subseteq \mathcal{U}(\mathbb{C}^d)$ of size $O(d^7/\epsilon^2)$ such that:*

1. N is an ϵ -pseudo-net with respect to spectral norm over unitaries $\mathcal{U}(\mathbb{C}^d)$.
2. Given index $i \in \{1, \dots, |N|\}$, the i 'th operator \tilde{U}_i in the net can be computed in time $O(d^2 \log^2(d^{5/2}/\epsilon))$. Here, by i 'th operator, we mean with respect to a fixed canonical ordering set by the construction of N .

The proof is given in Appendix A.

4 k -Orthogonality and the Traversal Lemma

The key technical tool for proving our hardness results is the Traversal Lemma (Lemma 4.2), which we state and prove in this section. In Sections 4.1.1 and 4.1.2, we then show that this lemma is tight up to a polynomial factor and give a further study into the notion of k -orthogonality, respectively. We begin by introducing the notions of k -orthogonal states and k -orthogonal subspaces.

Definition 4.1 (k -orthogonal states and subspaces). *For $k \geq 1$, a pair of states $|v\rangle, |w\rangle \in (\mathbb{C}^d)^{\otimes n}$ is k -orthogonal if for all k -qudit unitaries U , we have $\langle w|U|v\rangle = 0$. We call subspaces $S, T \subseteq (\mathbb{C}^d)^{\otimes n}$ k -orthogonal if any pair of vectors $|v\rangle \in S$ and $|w\rangle \in T$ are k -orthogonal.*

Let us comment on the structure of k -orthogonal states. First, k -orthogonality implies orthogonality, but not vice versa. For example, $|000\rangle$ and $|111\rangle$ are 2-orthogonal and hence orthogonal. In contrast, $|000\rangle$ and $|100\rangle$ are orthogonal but not k -orthogonal for any $k \geq 1$ (i.e. simply apply Pauli X to qubit 1 to map $|000\rangle$ to $|100\rangle$). Similarly, letting S and T denote the $+1$ eigenspaces of $I \otimes |000\rangle\langle 000|$ and $I \otimes |111\rangle\langle 111|$, respectively, we have that S and T are 2-orthogonal subspaces.

We now prove the Traversal Lemma, which says the following: For any two k -orthogonal subspaces S and T with $|v\rangle \in S$ and $|w\rangle \in T$, any sequence of m k -qudit unitaries mapping $|v\rangle$ to $|w\rangle$ must induce an evolution which has “large” overlap with the orthogonal complement of both S and T at some time step $i \in [m]$.

Lemma 4.2 (Traversal Lemma). *Let $S, T \subseteq (\mathbb{C}^d)^{\otimes n}$ be k -orthogonal subspaces. Fix arbitrary states $|v\rangle \in S$ and $|w\rangle \in T$, and consider a sequence of k -qudit unitaries $(U_i)_{i=1}^m$ such that*

$$\| |w\rangle - U_m \cdots U_1 |v\rangle \|_2 \leq \epsilon$$

for some $0 \leq \epsilon < 1/2$. Define $|v_i\rangle := U_i \cdots U_1 |v\rangle$ and $P := I - \Pi_S - \Pi_T$. Then, there exists an $i \in [m]$ such that

$$\langle v_i | P | v_i \rangle \geq \left(\frac{1 - 2\epsilon}{2m} \right)^2.$$

Proof. We give a proof by contradiction. Suppose that for all $i \in [m]$, the inner products satisfy $\langle v_i | P | v_i \rangle < \delta := [(1 - 2\epsilon)/(2m)]^2$. Consider the following thought experiment inspired by the quantum Zeno effect. Imagine that after each U_i is applied, we measure $|v_i\rangle$ using the projective measurement $(\Pi, I - \Pi)$ for $\Pi := I - P$, and postselect on obtaining outcome Π . Define the following two sequences:

- $|v'_i\rangle := \Pi |v_i\rangle$ for $i \in [m]$,
- $|v''_1\rangle := |v'_1\rangle$ and $|v''_i\rangle := \Pi U_i |v''_{i-1}\rangle$ for $i \in \{2, \dots, m\}$.

Note that $|v'_i\rangle$ and $|v''_i\rangle$ are not necessarily normalized.

To set up our contradiction, we first prove by induction on i that

$$\| |v_i\rangle\langle v_i| - |v''_i\rangle\langle v''_i| \|_{\text{tr}} < 2i\sqrt{\delta}. \quad (4)$$

For the base case $i = 1$, we have $|v''_1\rangle = |v'_1\rangle$. Then, since $\langle v_1 | P | v_1 \rangle < \delta$, we know that $\text{Tr}(\Pi |v_1\rangle\langle v_1|) > 1 - \delta$, and so the Gentle Measurement Lemma [Win99] (Lemma 2.3) yields

$$\| |v_1\rangle\langle v_1| - |v''_1\rangle\langle v''_1| \|_{\text{tr}} = \| |v_1\rangle\langle v_1| - |v'_1\rangle\langle v'_1| \|_{\text{tr}} < 2\sqrt{\delta}, \quad (5)$$

as required. For the inductive case, assume Equation (4) holds for $1 \leq i \leq j-1$. We prove it holds for $i = j$. Specifically,

$$\begin{aligned} \| |v_j\rangle\langle v_j| - |v''_j\rangle\langle v''_j| \|_{\text{tr}} &\leq \| |v_j\rangle\langle v_j| - |v'_j\rangle\langle v'_j| \|_{\text{tr}} + \| |v'_j\rangle\langle v'_j| - |v''_j\rangle\langle v''_j| \|_{\text{tr}} \\ &< 2\sqrt{\delta} + \| |v'_j\rangle\langle v'_j| - |v''_j\rangle\langle v''_j| \|_{\text{tr}} \\ &= 2\sqrt{\delta} + \left\| \Pi U_j (|v_{j-1}\rangle\langle v_{j-1}| - |v''_{j-1}\rangle\langle v''_{j-1}|) U_j^\dagger \Pi \right\|_{\text{tr}} \\ &\leq 2\sqrt{\delta} + \| |v_{j-1}\rangle\langle v_{j-1}| - |v''_{j-1}\rangle\langle v''_{j-1}| \|_{\text{tr}} \\ &< 2\sqrt{\delta} + 2(j-1)\sqrt{\delta_m} \\ &= 2j\sqrt{\delta}, \end{aligned} \quad (6)$$

where the first statement follows from the triangle inequality, the second from the Gentle Measurement Lemma, the fourth from the facts that the Schatten p -norms are invariant under isometries and that $\|ABC\|_p \leq \|A\|_\infty \|B\|_p \|C\|_\infty$ [Wat08], and the fifth from the induction hypothesis. This establishes Equality (4).

We thus have

$$\begin{aligned} \| |v''_m\rangle\langle v''_m| - |w\rangle\langle w| \|_{\text{tr}} &\leq \| |v''_m\rangle\langle v''_m| - |v_m\rangle\langle v_m| \|_{\text{tr}} + \| |v_m\rangle\langle v_m| - |w\rangle\langle w| \|_{\text{tr}} \\ &< 2m\sqrt{\delta} + 2\epsilon \\ &= 1, \end{aligned} \quad (7)$$

where we have used Equation (2) to bound

$$\| |v_m\rangle\langle v_m| - |w\rangle\langle w| \|_{\text{tr}} \leq 2 \| |v_m\rangle - |w\rangle \|_2 \leq 2\epsilon.$$

We are now ready to obtain the desired contradiction.

To do so, observe that since $|v\rangle \in S$, and since S and T are k -orthogonal subspaces, we have that for all $i \in [m]$, $|v''_i\rangle \in S$ (i.e., if S is 1-dimensional, this is the Zeno effect). Thus, we have $\langle v''_m | w \rangle = 0$, implying that

$$\| |v''_m\rangle\langle v''_m| - |w\rangle\langle w| \|_{\text{tr}} = 1 + \| |v''_m\rangle \|_2 \geq 1.$$

This contradicts Equation (7), as desired. \square

4.1 Tightness of the Traversal Lemma and properties of k -orthogonality

In the next two subsections, we discuss tightness of the Traversal Lemma and study the properties of k -orthogonality further.

4.1.1 On the tightness of the Traversal Lemma

We now ask whether the Traversal Lemma is tight in the following sense: In Lemma 4.2, the lower bound on $\langle v_i | P | v_i \rangle$ scales as $\Theta(1/m^2)$ (for m the number of unitaries and for fixed ϵ). This intuitively suggests that one can better “avoid” the subspace P projects onto if one uses a longer sequence of local unitaries. Is such behavior possible? Or can the lower bound in Lemma 4.2 be improved to a constant independent of m ? In this section, we show that a dependence on m in Lemma 4.2 is indeed necessary.

Theorem 4.3. *We assume the notation of Lemma 4.2. Fix any $0 < \Delta < 1/2$, and consider 2-orthogonal states $|v\rangle = |000\rangle$ and $|w\rangle = |111\rangle$, with $P := I - |v\rangle\langle v| - |w\rangle\langle w|$. Then, there exists a sequence of m 2-local unitary operations mapping $|v\rangle$ to $|w\rangle$ through intermediate states $|v_i\rangle$, each of which satisfy $\langle v_i | P | v_i \rangle \leq \Delta$, and where $m \in O(1/\Delta^2)$.*

The idea behind the proof is based on the following rough analogy: Suppose one wishes to map the point $(1, 1)$ (corresponding to $|000\rangle$) in the 2D Euclidean plane to $(-1, -1)$ (corresponding to $|111\rangle$) via a sequence of moves with the following two restrictions: (1) For each current point (x, y) , the next move must leave precisely one of x or y invariant (analogous to 2-local unitaries acting on a 3-qubit state), and (2) the Euclidean distance between (x, y) and the line through $(1, 1)$ and $(-1, -1)$ never exceeds Δ (analogous to the overlap with P not exceeding Δ). In other words, we wish to stay close to a diagonal line while making only horizontal and vertical moves. This can be achieved by making a sequence of “small” moves resembling a “staircase”. The smaller the size of each “step” in the staircase, the better we approximate the line, at the expense of requiring more moves (analogous to increasing the number of unitaries, m). Although the idea in this analogy is appealing in its simplicity, applying it to the setting of the Traversal Lemma is non-trivial, requiring a careful selection of 2-local unitary operations.

Proof of Theorem 4.3. Our high level approach is as follows. We first give a unitary U which is a sequence of two-qubit unitaries mapping $|000\rangle$ to $(|000\rangle + |111\rangle)/\sqrt{2}$. Given the technique behind U ’s construction, one can analogously obtain a unitary V which maps $(|000\rangle + |111\rangle)/\sqrt{2}$ to $|111\rangle$. It follows that $VU|000\rangle = |111\rangle$. It thus suffices to describe U , which is done in two steps. The first step consists of a pair of unitaries which transfer a small amount of amplitude from $|000\rangle$ to $|111\rangle$; applying this step repeatedly yields a state $|\psi\rangle$ “close” to $(|000\rangle + |111\rangle)/\sqrt{2}$. It is this iterative repetition which causes the overall number of unitaries m to scale as $\Omega(1/\Delta)$. Step 2 then maps $|\psi\rangle$ precisely onto $(|000\rangle + |111\rangle)/\sqrt{2}$. We now describe these steps.

Step 1: Iteratively make small steps towards $(|000\rangle + |111\rangle)/2$. Given any state of the form $\gamma_1|000\rangle + \gamma_2|111\rangle$ for real $\gamma_1 > \gamma_2$, we give a pair of two-qubit unitaries (U_1, U_2) which intuitively transfer a small amount of amplitude from $|000\rangle$ to $|111\rangle$.

We first apply a two-qubit unitary U_1 to qubits 1 and 2 with action $|00\rangle \mapsto \alpha|00\rangle + \beta|11\rangle$ and $|11\rangle \mapsto \beta|00\rangle - \alpha|11\rangle$ (for real α, β to be specified later), obtaining:

$$\begin{aligned} \gamma_1|000\rangle + \gamma_2|111\rangle &\mapsto \gamma_1\alpha|000\rangle + \gamma_1\beta|110\rangle + \gamma_2\beta|001\rangle - \gamma_2\alpha|111\rangle \\ &= |0\rangle(\gamma_1\alpha|00\rangle + \gamma_2\beta|01\rangle) + |1\rangle(\gamma_1\beta|10\rangle - \gamma_2\alpha|11\rangle). \end{aligned} \quad (8)$$

The overlap of this state with P is β^2 .

Next, apply a unitary U_2 on qubits 2 and 3 with action (omitting normalization for clarity) $(\gamma_1\alpha|00\rangle + \gamma_2\beta|01\rangle) \mapsto |00\rangle$ and $(\gamma_1\beta|10\rangle - \gamma_2\alpha|11\rangle) \mapsto |11\rangle$, obtaining:

$$\sqrt{\gamma_1^2\alpha^2 + \gamma_2^2\beta^2}|000\rangle + \sqrt{\gamma_1^2\beta^2 + \gamma_2^2\alpha^2}|111\rangle$$

which has 0 overlap with P . Setting $\beta = \sqrt{\Delta}$ ensures this process has at most Δ overlap with P at each intermediate step.

Let us now analyze the rate at which amplitude is transferred from $|000\rangle$ to $|111\rangle$ by this mapping. To do so, define

$$f(\gamma_1) := \sqrt{\gamma_1^2(1 - \Delta) + (1 - \gamma_1^2)\Delta} = \sqrt{(1 - 2\Delta)\gamma_1^2 + \Delta},$$

which is the new amplitude after the map induced by U_2U_1 is applied to input amplitude γ_1 . Note that $f(\gamma_1) \geq 0$ for all γ_1 and $f(\gamma_1)^2 > 1/2$ when $\gamma_1^2 > 1/2$.

We now quantify how *much* amplitude is transferred from $|000\rangle$ to $|111\rangle$ by this process. Suppose we iteratively apply U_2U_1 so long as $\gamma_1^2 \geq 1/2 + \zeta$ for some cutoff $\zeta \in (0, 1/2)$. Then, the difference $\gamma_1^2 - f(\gamma_1)^2$ satisfies

$$2\Delta\zeta \leq \gamma_1^2 - f(\gamma_1)^2 = \Delta(2\gamma_1^2 - 1) \leq \Delta.$$

In other words, each iterative step moves us at least $2\Delta\zeta$ and at most Δ towards our cutoff $1/2 + \zeta$, implying the number of iterations required to reach the cutoff scales between $\Omega(1/\Delta)$ and $O(1/\Delta\zeta)$. Setting $\zeta := 2\Delta/(1 + 2\Delta)$ (which lies in $(0, 1/2)$ when $\Delta \in (0, 1/2)$) ensures that the number of iterations is at most $O(1/\Delta^2)$, and sets up the next step of our transformation, which we now discuss.

Step 2: Map an “almost” equal superposition to an equal superposition. After $O(1/\Delta^2)$ iterations of Step 1, we arrive at a state of the form $\gamma_1|000\rangle + \gamma_2|111\rangle$ where $\gamma_1 \geq 0$ satisfies

$$\frac{1}{2} < \gamma_1^2 \leq \frac{1}{2} + \frac{2\Delta}{1 + 2\Delta}. \quad (9)$$

We seek a sequence of one and two-qubit unitaries which map this state to $(|000\rangle + |111\rangle)/\sqrt{2}$. To attain this, we instead equivalently give a sequence of unitaries which achieves the reverse mapping.

To begin, we apply a unitary to qubits 1 and 2 with action $|00\rangle \mapsto |00\rangle$ and $|11\rangle \mapsto (\beta|0\rangle + \alpha|1\rangle)|1\rangle$ for real parameters α, β to be specified later. This maps $(|000\rangle + |111\rangle)/\sqrt{2}$ to

$$\frac{1}{\sqrt{2}}|000\rangle + \frac{\beta}{\sqrt{2}}|011\rangle + \frac{\alpha}{\sqrt{2}}|111\rangle = \delta|0\rangle \left(\frac{1}{\sqrt{2}\delta}|00\rangle + \frac{\beta}{\sqrt{2}\delta}|11\rangle \right) + \frac{\alpha}{\sqrt{2}}|1\rangle|11\rangle, \quad (10)$$

where $\delta := \sqrt{(1 + \beta^2)/2}$. The overlap with P at this point is $\beta^2/2$.

Next, define a unitary with action:

$$\frac{1}{\sqrt{2}\delta}|00\rangle + \frac{\beta}{\sqrt{2}\delta}|11\rangle \mapsto |00\rangle, \quad \frac{\beta}{\sqrt{2}\delta}|00\rangle - \frac{1}{\sqrt{2}\delta}|11\rangle \mapsto |11\rangle, \quad |01\rangle \mapsto |01\rangle, \quad |10\rangle \mapsto |10\rangle.$$

Since this unitary is Hermitian, it follows that $|11\rangle$ is mapped to $\frac{\beta}{\sqrt{2}\delta}|00\rangle - \frac{1}{\sqrt{2}\delta}|11\rangle$; hence, applying this unitary to qubits 2 and 3 in Eqn. (10) yields:

$$\delta|000\rangle + \frac{\alpha\beta}{2\delta}|100\rangle - \frac{\alpha}{2\delta}|111\rangle = \left(\delta|00\rangle + \frac{\alpha\beta}{2\delta}|10\rangle \right) |0\rangle - \frac{\alpha}{2\delta}|11\rangle |1\rangle.$$

This state has overlap $\alpha^2\beta^2/(2\delta)^2 \leq \beta^2/2$ with P .

Finally, apply a unitary on qubits 1 and 2 which maps $|11\rangle$ to $-|11\rangle$ and (the normalized version of) $\delta|00\rangle + \frac{\alpha\beta}{2\delta}|10\rangle$ to $|00\rangle$, obtaining

$$\sqrt{1 - \frac{\alpha^2}{4\delta^2}}|000\rangle + \frac{\alpha}{2\delta}|111\rangle. \quad (11)$$

It remains to set β so as (1) to prevent overlap more than Δ with P , i.e. we require $\beta^2/2 \leq \Delta$, and (2) to ensure that the amplitude on $|000\rangle$ in Eqn. (11) is precisely γ_1 . Defining β implicitly via the equation

$$\sqrt{1 - \frac{\alpha^2}{4\delta^2}} = \sqrt{1 - \frac{1 - \beta^2}{2(1 + \beta^2)}} = \gamma_1$$

clearly satisfies the second of these requirements. Using the upper bound on γ_1^2 from Eqn. (9), it is straightforward to verify that the first requirement is also met. \square

4.1.2 Properties of k -orthogonality

We now study the properties of k -orthogonality further, and give an intuitive characterization of the notion (Lemma 4.5). We hope this may prove useful in possible independent applications of the concepts introduced in this section.

We begin with the following useful lemma.

Lemma 4.4. *For any $|v\rangle, |w\rangle \in (\mathbb{C}^d)^{\otimes n}$, $|v\rangle$ and $|w\rangle$ are k -orthogonal if and only if for all subsets of qudits $S \subseteq [n]$ of size at most k , we have $\text{Tr}_{[n]\setminus S}(|v\rangle\langle w|) = 0$.*

Proof. Assume first that $|v\rangle$ and $|w\rangle$ are k -orthogonal, and consider any $S \subseteq [n]$ with $|S| \leq k$. Then, we have

$$0 = \max_{U_S \in \text{U}((\mathbb{C}^d)^{\otimes |S|})} |\langle w| I \otimes U_S |v\rangle| = \max_{U_S \in \text{U}((\mathbb{C}^d)^{\otimes |S|})} |\langle U_S, \text{Tr}_{[n]\setminus S}(|v\rangle\langle w|) \rangle| = \|\text{Tr}_{[n]\setminus S}(|v\rangle\langle w|)\|_{\text{tr}},$$

where the second equality follows since $\text{Tr}((I_A \otimes C_B)D_{AB}) = \text{Tr}(C \text{Tr}_A(D_{AB}))$ for all linear operators C and D , and the third equality since $\|A\|_{\text{tr}} = \max_{U \in \text{U}(\mathcal{X})} |\text{Tr}(UA)|$ [Wat08] (intuitively, this holds since the optimal U rotates the set of left singular vectors of A into the set of right singular vectors of A). But now the claim follows, since $\|A\|_{\text{tr}} = 0$ if and only if $A = 0$. The converse direction proceeds analogously. \square

Lemma 4.5. *For any $|v\rangle, |w\rangle \in (\mathbb{C}^d)^{\otimes n}$, $|v\rangle$ and $|w\rangle$ are k -orthogonal if and only if for all subsets of qudits $S \subseteq [n]$ of size at most k , we have $(\text{Tr}_S |v\rangle\langle v|)(\text{Tr}_S |w\rangle\langle w|) = 0$.*

Proof. Assume first that $|v\rangle$ and $|w\rangle$ are k -orthogonal, and consider any $S \subseteq [n]$ with $|S| \leq k$. Let Y denote the register corresponding to $[n] \setminus S$. Then, suppose the Schmidt decompositions of $|v\rangle$ and $|w\rangle$ are $|v\rangle = \sum_i \alpha_i |a_i\rangle_S |b_i\rangle_Y$ and $|w\rangle = \sum_j \beta_j |c_j\rangle_S |d_j\rangle_Y$, respectively. Now, by Lemma 4.4, we have

$$0 = \text{Tr}_Y(|v\rangle\langle w|) = \sum_{ij} \alpha_i \beta_j \langle d_j | b_i \rangle |a_i\rangle \langle c_j|.$$

Since $\{|a_i\rangle\}$ and $\{|c_j\rangle\}$ are orthonormal sets, this implies that for any i and j , either $\alpha_i = 0$, $\beta_j = 0$, or $\langle d_j | b_i \rangle = 0$. Thus, letting

$$\rho := \text{Tr}_S(|v\rangle\langle v|) = \sum_i \alpha_i^2 |b_i\rangle\langle b_i| \quad \text{and} \quad \sigma := \text{Tr}_S(|w\rangle\langle w|) = \sum_j \beta_j^2 |d_j\rangle\langle d_j|,$$

we have $\text{Tr}(\rho\sigma) = \sum_{ij} \alpha_i^2 \beta_j^2 |\langle d_j | b_i \rangle|^2 = 0$, or equivalently, $\rho\sigma = 0$ since $\rho, \sigma \succeq 0$. The converse direction proceeds analogously. \square

Using Lemma 4.4, we can also easily show the following statement, which studies a notion reminiscent of positive maps versus *completely* positive maps.

Lemma 4.6. *For any $|v\rangle, |w\rangle \in (\mathbb{C}^d)^{\otimes n}$, $|v\rangle$ and $|w\rangle$ are k -orthogonal if and only if for any $|V\rangle, |W\rangle \in (\mathbb{C}^d)^{\otimes n'}$ for $n' \geq 0$, $|v\rangle |V\rangle$ and $|w\rangle |W\rangle$ are k -orthogonal.*

Proof. Assume first that $|v\rangle$ and $|w\rangle$ are k -orthogonal, and consider arbitrary $n' \geq 1$ and vectors $|V\rangle, |W\rangle \in (\mathbb{C}^d)^{\otimes n'}$. Let $S \subseteq [n]$ and $S' \subseteq [n']$ be such that $|S \cup S'| \leq k$. Then we have

$$\text{Tr}_{[n] \cup [n'] \setminus (S \cup S')}(|V\rangle |v\rangle \langle W| \langle w|) = \text{Tr}_{[n] \setminus S}(|V\rangle \langle W|) \otimes \text{Tr}_{[n'] \setminus S'}(|v\rangle \langle w|) = 0,$$

where the last equality holds since $|v\rangle$ and $|w\rangle$ are k -orthogonal and by Lemma 4.4. Thus, $|V\rangle |v\rangle$ and $|W\rangle |w\rangle$ are k -orthogonal. Since $|V\rangle$ and $|W\rangle$ are arbitrary, this direction of the claim holds. The converse statement is trivially true. \square

5 QCMA-completeness

In this section, we prove the following theorem.

Theorem 5.1. *There exists a polynomial p such that GSCON is QCMA-complete for $m \in O(p(n))$, $\Delta \in \Theta(1/m^5)$, $l = 2$, and $k \geq 5$, where n denotes the number of qubits H acts on.*

Intuitively, this says that GSCON is QCMA-complete when the unitaries U_i are at most 2-local, the number of unitaries scales polynomially, and the gap Δ scales inverse polynomially. To prove this theorem, we prove QCMA-hardness and containment in QCMA separately. We begin with QCMA-hardness.

5.1 QCMA-hardness

We now show that GSCON is QCMA-hard in the regime described below.

Lemma 5.2. *There exists a polynomial p such that GSCON is QCMA-hard for $m \in O(p(n))$, $\Delta \in O(1/m^5)$, $l = 2$, and $k \geq 5$, where n denotes the number of qubits H acts on.*

Proof. At a high level, our approach is as follows. Given a QCMA verification circuit V , let H' be the 3-local Hamiltonian output by Kempe and Regev’s circuit-to-Hamiltonian construction. Then, our aim is to construct another Hamiltonian H such that “traversing the ground space of H ” forces one to simulate the following protocol — starting with an initial state of all zeroes:

1. Apply a sequence of 2-qubit gates to prepare a ground state $|\psi_{H'}\rangle$ of H' .
2. Flip a first “GO” qubit to initiate a “check” that $|\psi_{H'}\rangle$ is indeed a ground state of H' .
3. Flip a second and third “GO” qubit to end the “check”.
4. Uncompute $|\psi_{H'}\rangle$ to obtain a target state which is all zeroes, except for the “GO” qubits, which are set to all ones.

Formally, let Π' be an instance of a QCMA problem with verification circuit V' acting on a classical proof register p and ancilla register a consisting of n_p and n_a qubits, respectively. Using standard error reduction via parallel repetition, we may assume without loss of generality that V' accepts (rejects) in the YES (NO) case with probability at least $p_{\text{accept}} \geq 1 - 2^{-\langle \Pi' \rangle}$ ($p_{\text{reject}} \geq 1 - 2^{-\langle \Pi' \rangle}$), where $\langle \Pi' \rangle$ denotes the encoding length of Π' .

Let V denote a new circuit which first measures the proof register in the computational basis, and then runs V' . Formally, V has the following properties: (1) V has $n_a + n_p$ ancilla qubits initialized to all zeroes, (2) in time step $i \in [n_p]$, V applies a CNOT gate with the i 'th proof qubit as control and ancilla qubit $n_a + i$ as target, and (3) starting at time step $n_p + 1$, V simulates V' while acting on register p and the first n_a qubits of a . A straightforward argument shows that V accepts a proof if and only if V' does. Moreover, unlike V' , the principle of deferred measurement [NC00] yields that V is sound against a cheating prover which does not send a classical string x as a proof.

Next, we define our Hamiltonian H based on V . Let H' denote the 3-local Hamiltonian obtained from V using Kempe and Regev's circuit-to-Hamiltonian construction [KR03]. Then, we define H to act on a *Hamiltonian* register denoted h and *GO* register denoted G . Specifically,

$$H \in \text{Herm} \left((\mathbb{C}^2)^{\otimes (2n_p + n_a + n_c)} \otimes (\mathbb{C}^2)^{\otimes 3} \right),$$

where n_c denotes the polynomial number of qubits used for the clock register of H' , and

$$H := H'_h \otimes P_G \quad \text{for} \quad P := I - |000\rangle\langle 000| - |111\rangle\langle 111|. \quad (12)$$

Noting that P can be written 2-locally as

$$\begin{aligned} P = & \frac{1}{2} (|01\rangle\langle 01| \otimes I + |10\rangle\langle 10| \otimes I + I \otimes |01\rangle\langle 01| + I \otimes |10\rangle\langle 10| + \\ & |1\rangle\langle 1| \otimes I \otimes |0\rangle\langle 0| + |0\rangle\langle 0| \otimes I \otimes |1\rangle\langle 1|), \end{aligned}$$

we have that H is 5-local. We define our initial and final states as

$$|\psi\rangle := |0\rangle^{\otimes (2n_p + n_a + n_c)} |0\rangle^{\otimes 3} \quad \text{and} \quad |\phi\rangle := |0\rangle^{\otimes (2n_p + n_a + n_c)} |1\rangle^{\otimes 3}. \quad (13)$$

Finally, letting W denote a unitary circuit of size $|W|$ which prepares the history state of H given classical proof x , define $m := 2(n_p + |W| + 1)$. Note that m is polynomial in the input size, since for any YES instance Π , V' accepts a *classical* proof, and hence the history state for H' can be prepared in polynomial time. (This observation was also made in [WJB03].) Set $\eta_3 = 0$, $\eta_4 = 1/4$, $\eta_1 = \alpha$, and $\eta_2 = \beta/(16m^2)$, where α and β come from Lemma 2.4. Thus, $\eta_1 \in O(2^{-\langle \Pi' \rangle})$ and $\eta_2 \in \Omega(1/m^5) \in \Omega(1/\text{poly}(\langle \Pi' \rangle))$ (where we have used the facts that $m \geq L$ for L the number of gates in circuit V and $m \in \text{poly}(\langle \Pi' \rangle)$). Choose $\Delta \in O(1/m^5)$ and set $l = 2$. Observe that $\Pi = (H, \eta_1, \eta_2, \eta_3, \eta_4, \Delta, l, m, |\psi\rangle, |\phi\rangle)$ is a valid instance of GSCON which can be computed in polynomial time given $\Pi' = (V')$, as desired.

We now show correctness. Suppose there exists a proof $x \in \{0, 1\}^{n_p}$ accepted by V . We demonstrate a sequence $(U_i)_{i=1}^m$ of 2-qubit unitaries mapping $|\psi\rangle$ to $|\phi\rangle$ through the ground space of H . First, note that $|\psi\rangle$ and $|\phi\rangle$ are in the null space of H , and hence $\langle \psi | H | \psi \rangle \leq \eta_1$ and $\langle \phi | H | \phi \rangle \leq \eta_1$, as required. Next, recall in Kempe and Regev's construction that the Hamiltonian register h is itself composed of three sub-registers h_1 , h_2 , and h_3 , corresponding to the *proof*, *ancilla*, and *clock* registers for H , respectively. The desired sequence $(U_i)_{i=1}^m$ is then given as follows:

1. Apply Pauli X gates to h_1 to prepare classical proof x , i.e., map $|0\rangle^{\otimes n_p}$ to $|x\rangle$.
2. Apply W to h to prepare the history state $|\text{hist}_x\rangle$ of H' .

3. Apply $(X \otimes X \otimes I)_G$ to “initiate” checking of $|\text{hist}_x\rangle$.
4. Apply $(I \otimes I \otimes X)_G$ to “complete” checking of $|\text{hist}_x\rangle$.
5. Apply W^\dagger to h to uncompute $|\text{hist}_x\rangle$.
6. Apply X gates to h_1 to map the initial proof $|x\rangle$ back to $|0\rangle^{\otimes n_p}$.

Note first that the length of the sequence above is at most $2(n_p + |W| + 1)$ gates, as desired. Second, the final state is equal to $|\phi\rangle$, and every intermediate state is in the null space of H except for possibly after Step 3. As for after Step 3, let $|a_3\rangle$ denote our state at this point. Then, since a valid history state $|\text{hist}_x\rangle$ obtains energy $\langle \text{hist}_x | H' | \text{hist}_x \rangle \leq \alpha$, we have $\langle a_3 | H | a_3 \rangle \leq \alpha = \eta_1$, as desired. Thus, if Π' is a YES instance, then Π is a YES instance of GSCON.

Conversely, suppose Π' is a NO instance, i.e., for all $x \in \{0, 1\}^{n_p}$, V rejects with high probability. Then, by Lemma 2.4, the smallest eigenvalue of H' is at least β . Now, let S and T denote the $+1$ eigenspaces of projections $I_h \otimes |000\rangle\langle 000|_G$ and $I_h \otimes |111\rangle\langle 111|_G$, respectively. Observe that S and T are 2-orthogonal subspaces, and that $|\psi\rangle \in S$ and $|\phi\rangle \in T$. Thus, for any sequence of two-qubit unitaries $(U_i)_{i=1}^m$, either $\| |\psi_m\rangle - |\phi\rangle \|_2 \geq 1/4 = \eta_4$ (in which case we have a NO instance of GSCON and we are done), or we can apply the Traversal Lemma (Lemma 4.2) with $\epsilon = 1/4$ to conclude that there exists an $i \in [m]$ such that

$$\langle \psi_i | P' | \psi_i \rangle \geq \left(\frac{1}{4m} \right)^2 = \frac{\eta_2}{\beta},$$

where we define $|\psi_i\rangle := U_i \cdots U_1 |\psi\rangle$ and $P' = I - \Pi_S - \Pi_T$. Note that we can write P' as $I_h \otimes P$. We conclude that

$$\langle \psi_i | H | \psi_i \rangle = \langle \psi_i | H' \otimes P | \psi_i \rangle \geq \beta \langle \psi_i | I_h \otimes P | \psi_i \rangle = \beta \langle \psi_i | P' | \psi_i \rangle \geq \eta_2,$$

where the first inequality follows since $H' \succeq \beta I$. □

Remark. There is no loss of generality in restricting ourselves to 2-qubit unitaries in the proof above. Specifically, the same proof applies almost identically if we instead allow p -qubit unitaries for any constant $p \geq 2$ by changing Equation (13) to

$$|\psi\rangle := |0\rangle^{\otimes(2n_p+n_a+n_c)} |0\rangle^{\otimes(p+1)} \quad \text{and} \quad |\phi\rangle := |0\rangle^{\otimes(2n_p+n_a+n_c)} |1\rangle^{\otimes(p+1)},$$

i.e., the GO register consists more generally of $p+1$ qubits. Note that the Traversal Lemma still applies in this more general setting, and second, the projector P onto the GO register can be represented as a 2-local Hamiltonian regardless of the value of p , implying we still have $k = 5$.

Remark. In the proof of Theorem 5.2, we used Kempe and Regev’s 3-local circuit-to-Hamiltonian construction. One might ask whether one of the known 2-local constructions based on perturbation theory gadgets may instead be applied to reduce the locality of H further. The main issue in doing so is that here we require the ability to construct the ground state efficiently. In other words, the perturbation theory reduction should ideally produce a ground state whose structure is similar to the history state. Now, Oliveira and Terhal [OT08] have in fact proven such a perturbation theory result in which the resulting 2-local Hamiltonian’s ground space approximates the starting Hamiltonian’s ground space. However, we require a stronger statement than this. To explain, let H denote a k -local Hamiltonian and H' the 2-local Hamiltonian resulting from the construction

in [OT08]. Then, our proof requires a statement of the form²: If $\langle v | H | v \rangle \leq a$, then $\langle v | H' | v \rangle \leq a$, and if $\langle v | H | v \rangle \geq b$, then $\langle v | H' | v \rangle \geq b$. Unfortunately, as far as we are aware, it seems the first of these conditions can be violated for the gadgets presented in [OT08]. Intuitively, what is happening here is that although $\langle v | H | v \rangle \leq a$ (i.e. the expectation is “small”), it may be that $|v\rangle$ does not fully lie in the ground space of H , but rather has some small overlap with a higher energy subspace S . If this higher energy space S is then penalized strongly in H' , then $\langle v | H' | v \rangle$ can be large.

5.2 Containment in QCMA

We now show that GSCON with 2-local unitaries U_i is in QCMA so long as the gap Δ scales inverse polynomially and the number of unitaries m scales polynomially with the input size.

Lemma 5.3. *For any nonnegative constants c_1 and c_2 , GSCON is in QCMA for $\Delta \geq 1/n^{c_1}$, $m \leq n^{c_2}$, $l = 2$, and $k \in O(\log n)$, where n denotes the number of qubits H acts on.*

Proof. Let $\Pi = (H, \eta_1, \eta_2, \eta_3, \eta_4, \Delta, l, m, |\psi\rangle, |\phi\rangle)$ be an instance of GSCON. The proof system is given below. Steps 4 and 5 follow standard ideas; thus, we simply sketch them here. Let L denote the number of local terms in H .

Algorithm 1. (QCMA proof system for GSCON)

1. *The prover sends a sequence $(\tilde{U}_i)_{i=1}^m \subseteq \mathcal{L}(\mathbb{C}^2 \otimes \mathbb{C}^2)$ of matrices from the ϵ -pseudo-net of Lemma 3.3, for $\epsilon := \Delta/16mL$.*
2. *(Unitary check) The verifier runs algorithm C from Lemma 3.3 on each \tilde{U}_i , and rejects if C rejects.*
3. *(Rounding step) The verifier uses algorithm R from Lemma 3.3 to construct a sequence $(V_i)_{i=1}^m \subseteq \mathcal{U}(\mathbb{C}^2 \otimes \mathbb{C}^2)$ such that for all $i \in [m]$, $\|\tilde{U}_i - V_i\|_\infty \leq \epsilon$.*
4. *(Low energy check) Define $|\psi_t\rangle := V_t \cdots V_1 |\psi\rangle$. For all $t \in [m]$, the verifier prepares state $|\psi_t\rangle$ a polynomial number of times, and runs Kitaev’s phase estimation procedure [KSV02] to estimate $\langle \psi_t | H | \psi_t \rangle$ within inverse polynomial accuracy. The verifier rejects if this estimate is larger than $\eta_1 + \epsilon$.*
5. *(Close to target state check) The verifier performs the SWAP test [BCWdW01] between $|\psi_m\rangle$ and $|\phi\rangle$ polynomially many times to estimate $\| |\psi_m\rangle - |\phi\rangle \|_2$ within inverse polynomial accuracy. The verifier rejects if this estimate is larger than $\eta_3 + \epsilon$.*
6. *The verifier accepts.*

The verifier’s action is clearly implementable by a polynomial size quantum circuit.

We now show correctness. Let N denote the ϵ -pseudo-net over 2-qubit unitaries from Lemma 3.3 (i.e., $d = 4$ in Lemma 3.3), for ϵ as chosen above. Suppose now that Π is a YES instance, i.e., there exists a sequence of 2-qubit unitaries $(U_i)_{i=1}^m$ mapping $|\psi\rangle$ to $|\phi\rangle$ through the ground space of H . Then, in Step 1, the prover sends sequence $(\tilde{U}_i)_{i=1}^m \in N^{\times m}$ such that $\|U_i - \tilde{U}_i\|_\infty \leq \epsilon$ for $i \in [m]$. By Definition 3.2 and Lemma 3.3, Step 2 will pass and the conditions of Step 3 will be met with certainty.

²Note that in [OT08], H and H' live in different spaces, so our statement here should not be read literally. Rather, it is intended to give a flavor of the ideal behavior we would like the perturbation theory reduction to obey, without getting into finer details in our discussion here.

Next, we claim that for all $t \in [m]$, $\|U_t \cdots U_1 - V_t \cdots V_1\|_\infty \leq 2\epsilon t$. To see this, we first bound

$$\|U_t \cdots U_1 - V_t \cdots V_1\|_\infty \leq \|U_t \cdots U_1 - \tilde{U}_t \cdots \tilde{U}_1\|_\infty + \|\tilde{U}_t \cdots \tilde{U}_1 - V_t \cdots V_1\|_\infty$$

and use the fact [NC00] that for any two quantum circuits $U = U_t \cdots U_1$ and $V = V_t \cdots V_1$ satisfying $\|U_j - V_j\|_\infty \leq \epsilon$, we have $\|U - V\|_\infty \leq \sum_{i=1}^t \|U_i - V_i\|_\infty$. Defining $|u_t\rangle := U_t \cdots U_1 |\psi\rangle$ and recalling that $|\psi_t\rangle := V_t \cdots V_1 |\psi\rangle$, it follows that for all $t \in [m]$, $\| |u_t\rangle - |\psi_t\rangle \|_2 \leq 2\epsilon m$. Thus,

$$|\text{Tr}(H |u_t\rangle\langle u_t|) - \text{Tr}(H |\psi_t\rangle\langle \psi_t|)| \leq \|H\|_\infty \| |u_t\rangle\langle u_t| - |\psi_t\rangle\langle \psi_t| \|_{\text{tr}} \leq 4\epsilon m L,$$

where recall L denotes the number of local terms in H , the first inequality follows from Hölder's inequality, and the second by Equation (2).

Since we chose $\epsilon = \Delta/16mL$, we have $(\eta_2 - 4\epsilon m L) - (\eta_1 + 4\epsilon m L) \geq \Delta/2$ and we also have $(\eta_4 - 2\epsilon m) - (\eta_3 + 2\epsilon m) \geq \Delta/2$, i.e., the error incurred by using our net N shifts the thresholds which Steps 4 and 5 must distinguish between by at most $\Delta/4$ each, leaving gaps of size $\Delta/2$. But $\Delta/2$ is inverse polynomially large; thus, with high probability (i.e., inverse exponentially close to 1), Steps 4 and 5 do not reject. We conclude that with high probability, the verifier accepts, as desired.

Conversely, suppose we have a NO instance. Then, either the verifier rejects in Step 2, or it runs Step 3 to “round” the prover’s provided matrices into a sequence of unitaries $(V_i)_{i=1}^m$. But by the NO conditions of GSCON, we know that for our choice of ϵ , either Step 4 or Step 5 must now reject with high probability (i.e., inverse exponentially close to 1). \square

6 PSPACE-completeness

In this section, we show the following theorem.

Theorem 6.1. *GSCON is PSPACE-complete for $m = 2^n$, $\Delta = 2^{-(2n+4)}$, $l = 1$, $k = 3$, where n denotes the number of qubits H acts on.*

Intuitively, this says that GSCON is PSPACE-complete when the unitaries are 1-local, the number of unitaries scales exponentially, and the gap Δ scales inverse exponentially. To show this, we prove PSPACE-hardness and containment in PSPACE separately. We begin with PSPACE-hardness.

6.1 PSPACE-hardness

We now show PSPACE-hardness of GSCON for the case of exponentially many 1-local unitaries and exponentially small gap Δ .

Lemma 6.2. *GSCON is PSPACE-hard for $k = 3$, $\eta_1 = \eta_3 = 0$, $\eta_2 = 2^{-(2n+4)}$, $\eta_4 = 1/4$, $\Delta = 2^{-(2n+4)}$, $l = 1$, and $m = 2^n$, where n denotes the number of qubits H acts on.*

Proof. We show a polynomial-time many-one or Karp reduction from s,t-CONN (which by Theorem 2.6 is PSPACE-complete) to GSCON. Specifically, let $\Pi = (\phi, x, y)$ be an instance of s,t-CONN for 3-CNF ϕ . The main idea is to embed ϕ trivially into a 3-local Hamiltonian H as follows. For each clause c_i of ϕ , we define a local Hamiltonian constraint H_i to penalize the unique 3-bit “bad” assignment to c_i , i.e., $H_i := |z_i\rangle\langle z_i|$ for $c_i(z_i) = 0$. Setting our parameters as in the theorem statement, we thus obtain an instance $\Pi' = (H := \sum_i H_i, \eta_1, \eta_2, \eta_3, \eta_4, \Delta, l, m, U_x, U_y)$ of GSCON, where $U_x |0 \cdots 0\rangle = |x\rangle$ and $U_y |0 \cdots 0\rangle = |y\rangle$ for the strings x and y , respectively, from the s,t-CONN instance. Now, given strings $x, y \in \{0, 1\}^n$, it is trivial that if Π is a YES instance

of s,t-CONN, then Π' is a YES instance of GSCON: Namely, simulate local bit flips on strings by Pauli X gates to map $|x\rangle$ to $|y\rangle$ while staying in the null space of H . Note that since there are at most 2^n distinct strings on n bits, at most $m = 2^n$ Pauli X gates suffice to map $|x\rangle$ to $|y\rangle$.

Conversely, suppose Π is a NO instance of s,t-CONN. Let S denote the subspace corresponding to the span of all states $|z\rangle$ such that z can be obtained via a sequence of bit flips from x , where each string in the sequence is a satisfying assignment to ϕ . Let T denote the span of all remaining satisfying assignments. Note that $|x\rangle \in S$, $|y\rangle \in T$. Also, the Hamming distance from any computational basis state in S to computational basis state in T is at least 2; thus, S and T are 1-orthogonal subspaces. From the Traversal Lemma (Lemma 4.2), we know for any sequence of one-qubit unitaries $(U_i)_{i=1}^m$ that either $\|\psi_m - |\phi\rangle\|_2 \geq \eta_4 = 1/4$, or there exists an $i \in [m]$ such that $\langle \psi_i | P' | \psi_i \rangle \geq (1/(4m))^2 = 2^{-(2n+4)}$, where we again define $|\psi_i\rangle := U_i \cdots U_1 |\psi\rangle$ and $P' = I - \Pi_S - \Pi_T$. Thus, if it were the case that $H \succeq P'$, then

$$\langle \psi_i | H | \psi_i \rangle \geq \langle \psi_i | P' | \psi_i \rangle \geq \frac{1}{2^{2n+4}} = \eta_2,$$

as desired. To see that indeed $H \succeq P'$, note that H and P' are diagonal matrices with non-negative integer entries satisfying for $z \in \{0, 1\}^n$:

$$(\langle z | H | z \rangle = 0 \iff \langle z | P' | z \rangle = 0) \quad \text{and} \quad (\langle z | H | z \rangle \geq 1 \iff \langle z | P' | z \rangle = 1).$$

This concludes the proof. □

6.2 Containment in PSPACE

We now show that GSCON is in PSPACE for exponentially many 1-local unitaries U_i and inverse exponential gap Δ .

Lemma 6.3. *For all nonnegative constants c_1 and c_2 , GSCON with $l = 1$ is in PSPACE for $m \leq 2^{n^{c_1}}$ and $\Delta \geq 1/2^{n^{c_2}}$, where n denotes the number of qubits H acts on.*

Proof. We give a non-deterministic polynomial space algorithm for GSCON, and subsequently apply Savitch's theorem [Sav70] to obtain a PSPACE algorithm. Specifically, given a GSCON instance $\Pi = (H, \eta_1, \eta_2, \eta_3, \eta_4, \Delta, l, m, |\psi\rangle, |\phi\rangle)$, our non-deterministic algorithm proceeds as follows. Let L denote the number of local terms in H , and let N denote the ϵ -net for single qubit unitaries from Lemma 3.1 for $\epsilon := \Delta/8L(2(m-1)+1)$. Then our algorithm is given by (explanation to follow):

Algorithm 2. *(Polynomial space algorithm for GSCON)*

1. If $\|\psi\rangle - |\phi\rangle\|_2 \leq \eta_3$, accept.
2. For $i \in \{0, \dots, m\}$, define $V_i := V_{i,1} \otimes \cdots \otimes V_{i,n}$ (for operators $V_{i,j}$ to be defined in iteration i).
3. Set $V_{0,j} := I$ for all $j \in [n]$, i.e., $V_0 := I$.
4. For $i = 1$ to m , do:
 - (a) Non-deterministically guess a unitary $B_i \in U(\mathbb{C}^2)$ from N , where B_i acts on some qubit $q \in [n]$ chosen non-deterministically.

- (b) For $j \neq q$, set $V_{ij} := V_{i-1,j}$. Set $V'_{iq} := B_i V_{i-1,q}$.
- (c) Set $V_{iq} := \text{round}(V'_{iq})$, where $\text{round}(A)$ straightforwardly maps unitary A to a net element $A' \in N$ such that $\|A - A'\|_\infty \leq \epsilon$.
- (d) (Energy Test) If $\langle \psi | V_i^\dagger H V_i | \psi \rangle \geq \eta_1 + \Delta/3$, exit loop.
- (e) (Proximity Test) If $\|V_i |\psi\rangle - |\phi\rangle\|_2 \leq \eta_3 + \Delta/4$, accept.

5. Reject.

The intuition behind the algorithm is as follows. Ideally, we would like to run the following algorithm: At each step, non-deterministically guess a unitary $U \in U(\mathbb{C}^d)$, apply U to the state computed in the previous step, and check whether the new state has high energy (Step 4(d)), or is close to the target state (Step 4(e)). Note that at a high level, this is possible in PSPACE because each unitary acts on a single qubit; thus, it suffices to keep track of the *cumulative* single-qubit unitary applied to each qubit after each step (Step 4(b)), as opposed to keeping a history of all m (i.e. exponentially many) unitaries guessed in Step 4. In particular, this implies the overall unitary V_i in each iteration has a *succinct* description (i.e., of tensor product form). There are, however, two subtle issues with this approach. The first is that the space of unitaries is continuous; thus, in iteration i , our algorithm non-deterministically chooses a unitary B_i from N instead (Step 4(a)). The second issue is that m is exponentially large — thus, multiplying all B_i which act on a qubit j can result in an operator whose entries require an exponential number of bits of precision. To prevent this, in each iteration, Step 4(c) “rounds” the product $B_i V_{i-1,q}$ back to an operator in our net.

We now justify why the algorithm runs in polynomial space. Since each V_i can be described using a polynomial number of bits, Step 4(a) can be carried out by a Turing machine whose configurations each require at most polynomially many bits to specify. For Step 4(c), since ϵ is inverse exponential in our setting, Lemma 3.1 implies $|N|$ scales exponentially; thus, Step 4(c) can be achieved in polynomial space via a brute force search over all indices i of operators in the net via Lemma 3.1. Steps 4(d) and 4(e) can be completed in polynomial space using the standard approach of recomputing any values needed on-the-fly when determining (say) an inner product of exponentially large vectors specified by polynomial-size quantum circuits. We conclude that the algorithm runs in polynomial space.

We now justify correctness. Suppose first that there exists a sequence of 1-local unitaries $(\hat{U}_i)_{i=1}^m$ satisfying the conditions of a YES instance of GSCON. For convenience, define the global unitary after step i as $U_i := U_{i,1} \otimes \cdots \otimes U_{i,n}$. We prove by induction on i that for all $i \in [m]$,

$$\|U_i - V_i\|_\infty \leq (2(i-1) + 1)\epsilon. \quad (14)$$

For the base case $i = 1$, we have $\|U_1 - V_1\|_\infty \leq \epsilon$ since $\|A \otimes B\|_\infty = \|A\|_\infty \|B\|_\infty$ (recall U_1 and V_1 act non-trivially only on a single qubit) and by Lemma 3.1. Thus, the base case holds. For the inductive step, assume the claim is true for iterations 1 through $i-1$. We prove it for iteration i . Specifically,

$$\begin{aligned} \|U_i - V_i\|_\infty &= \left\| \hat{U}_i U_{i-1} - \text{round}(B_i V_{i-1}) \right\|_\infty \\ &\leq \left\| \hat{U}_i U_{i-1} - B_i V_{i-1} \right\|_\infty + \|B_i V_{i-1} - \text{round}(B_i V_{i-1})\|_\infty \\ &\leq \left\| \hat{U}_i - B_i \right\|_\infty + \|U_{i-1} - V_{i-1}\|_\infty + \epsilon \\ &\leq \epsilon + (2(i-2) + 1)\epsilon + \epsilon \\ &= (2(i-1) + 1)\epsilon, \end{aligned}$$

where the first inequality follows from the triangle inequality, the second inequality from the fact that $\|AB - CD\|_\infty \leq \|A - C\|_\infty + \|B - D\|_\infty$ for unitaries A, B, C, D and by Lemma 3.1, and the third inequality from Lemma 3.1 and the induction hypothesis. This completes our proof of Equation (14).

We conclude that in any iteration $i \in [m]$, we have $\|U_i - V_i\|_\infty \leq (2(i-1) + 1)\epsilon$, and hence $\|U_i|\psi\rangle - V_i|\psi\rangle\|_2 \leq (2(i-1) + 1)\epsilon$. Recalling that L is the number of local terms in H , this yields

$$\begin{aligned} \left| \langle \psi | U_i^\dagger H U_i | \psi \rangle - \langle \psi | V_i^\dagger H V_i | \psi \rangle \right| &\leq \|H\|_\infty \left\| U_i | \psi \rangle \langle \psi | U_i^\dagger - V_i | \psi \rangle \langle \psi | V_i^\dagger \right\|_{\text{tr}} \\ &\leq 2L \|U_i | \psi \rangle - V_i | \psi \rangle\|_2 \\ &\leq 2L(2(i-1) + 1)\epsilon \\ &\leq \frac{\Delta}{4}, \end{aligned} \tag{15}$$

where the first inequality follows from Hölder's inequality, and the second from Equation (2). In addition, since in a YES instance $\|U_m |\psi\rangle - |\phi\rangle\|_2 \leq \eta_3$, by the triangle inequality we have

$$\|V_m |\psi\rangle - |\phi\rangle\|_2 \leq \|V_m |\psi\rangle - U_m |\psi\rangle\|_2 + \|U_m |\psi\rangle - |\phi\rangle\|_2 \leq (2(m-1) + 1)\epsilon + \eta_3 \leq \frac{\Delta}{4} + \eta_3. \tag{16}$$

By Equations (15) and (16), we conclude that for a YES instance of GSCON, Step 4(d) of our algorithm will never cause an exit from the loop, and Step 4(e) will accept in some iteration. An analogous argument shows that for any NO instance, either the algorithm exits the loop in Step 4(d) or never passes the check in Step 4(e), implying the algorithm rejects, as desired. \square

7 NEXP-completeness

In this section, we define a succinct version of GSCON, and show that it is NEXP-complete. As the proof techniques used here are essentially the same as in Sections 5 (QCMA-completeness) and 6 (PSPACE-completeness), for brevity we give only proof sketches.

We begin by defining succinct or *oracle* notions of a local Hamiltonian and quantum product states, in analogy with an oracle 3-CNF formula and oracle truth assignment [BR04].

Definition 7.1 (Oracle k -local Hamiltonian). *Let $H = \sum_{i=1}^{2^r} H_i$ be a k -local Hamiltonian acting on 2^n qubits with 2^r clauses. An oracle local Hamiltonian is a classical circuit C_H which, given index $i \in \{0, 1\}^r$ as input, outputs a classical description of constraint H_i (i.e. outputs a $2^k \times 2^k$ matrix), along with the indices of the k qubits on which H_i acts.*

Definition 7.2 (Oracle quantum product state). *Let $|\psi\rangle$ be a tensor product state on 2^n qubits such that $|\psi\rangle = |\psi_1\rangle \otimes |\psi_2\rangle \otimes \cdots \otimes |\psi_{2^n}\rangle$. An oracle quantum product state is a classical circuit C_ψ which, given index $i \in \{0, 1\}^n$ as input, outputs a classical description of $|\psi_i\rangle$.*

Using these two definitions, we can now define the succinct version of GSCON.

Definition 7.3 (SUCCINCT GSCON($H, k, \eta_1, \eta_2, \eta_3, \eta_4, \Delta, l, m, U_\psi, U_\phi$)). *SUCCINCT GSCON is defined identically to GSCON, except the Hamiltonian H is an oracle Hamiltonian and the initial states $|\psi\rangle$ and $|\phi\rangle$ are oracle quantum product states.*

In this section, we show the following theorem.

Theorem 7.4. *SUCCINCT GSCON is NEXP-complete for $m \in O(2^n)$, $\Delta \in \Theta(1/m^2)$, $l = 1$, and $k \geq 5$, where 2^n is the number of qubits H acts on.*

Intuitively, this says that the succinct version of GSCON in which (1) the number of unitaries scales linearly in the number of qubits (but exponentially in the input size) and (2) each unitary is 1-local is NEXP-complete.

7.1 NEXP-hardness

We now show NEXP-hardness of SUCCINCT GSCON.

Lemma 7.5. *SUCCINCT GSCON is NEXP-hard for $m \in O(2^n)$, $\Delta \in O(1/m^2)$, $l = 1$, and $k \geq 5$, where 2^n is the number of qubits H acts on.*

Proof. We sketch a polynomial-time many-one or Karp reduction from the NEXP-complete problem ORACLE 3SAT (see, e.g. [BR04]) to SUCCINCT GSCON. Specifically, in an ORACLE 3SAT instance, one is given as input an oracle 3-CNF formula η consisting of 2^n variables and 2^r clauses; η can be thought of as a circuit C_η which, given index $i \in \{0, 1\}^m$, outputs the i 'th clause and the indices of the variables on which the i 'th clause acts.

Our approach is as follows: We embed the oracle 3-CNF formula into an oracle 3-local Hamiltonian in the trivial way, and subsequently combine this with the construction of Lemma 5.2 (QCMA-hardness). Specifically, our oracle Hamiltonian C_H acts as follows: Given index i , it runs C_η on i to obtain the i 'th clause c_i . It then converts this to a diagonal Hamiltonian constraint H_i (for example, clause $x_1 \vee x_2 \vee x_3$ is mapped to the diagonal operator $\text{Diag}(1, 0, 0, 0, 0, 0, 0, 0)$), and returns constraint $H_i \otimes P$ for $P := I - |00\rangle\langle 00| - |11\rangle\langle 11|$. (Note that here P needs only act on 2 qubits since the unitaries U_i are 1-local.) The initial and final states are oracle quantum product states C_ψ and C_ϕ representing $|\psi\rangle := |0\rangle^{\otimes 2^n} |0\rangle^{\otimes 2}$ and $|\phi\rangle := |0\rangle^{\otimes 2^n} |1\rangle^{\otimes 2}$, respectively. (Clearly, C_ψ and C_ϕ have size $\text{poly}(n)$.) Set $\eta_1 := 0$, $\eta_2 := 1/(16m^2)$, $\eta_3 := 0$, $\eta_4 := 1/4$, $\Delta \in O(1/m^2)$, and $l = 1$. This concludes the construction of our SUCCINCT GSCON instance.

To show correctness, for a YES instance, we proceed analogously to Lemma 5.2, except now there is no history state to prepare; in particular, the sequence of m unitaries is given by:

1. Apply Pauli X gates to h_1 to prepare satisfying assignment x for η , i.e. map $|0\rangle^{\otimes 2^n}$ to $|x\rangle$.
2. Apply $(X \otimes I)_G$ to “initiate” checking of $|x\rangle$.
3. Apply $(I \otimes X)_G$ to “complete” checking of $|x\rangle$.
4. Apply X gates to h_1 to map the initial proof $|x\rangle$ back to $|0\rangle^{\otimes 2^n}$.

Clearly, this process requires at most $m = 2^{n+1} + 2$ single-qubit unitaries, as desired. The analysis for a NO instance proceeds essentially identically to Lemma 5.2; one need only replace β by 1. The reason this works is because $H := \sum_i H_i \succeq I$ since it is a sum of diagonal projections and there does not exist a classical string z such that $\langle z | H | z \rangle = 0$. \square

7.2 Containment in NEXP

We now show containment of SUCCINCT GSCON in NEXP.

Lemma 7.6. *SUCCINCT GSCON with $l = 1$ is in NEXP for $m \leq \text{poly}(2^n)$ and $\Delta \geq 1/\text{poly}(2^n)$, where 2^n is the number of qubits H acts on.*

Proof. The proof is essentially identical to that of Lemma 6.3 (containment in PSPACE), i.e. the verifier runs Algorithm 2. As the Hamiltonian involved now acts on exponentially many qubits, a few remarks regarding the implementation of Algorithm 2 are in order:

- The initial state $|\psi\rangle$, final state $|\phi\rangle$, and intermediate states $V_i|\psi\rangle$ are product states. Hence, the Energy Test (Step 4(d)) and Proximity Test (Step 4(e)) can be carried out in exponential time. For example, suppose for the former that we wish to estimate $\langle\psi|V_i^\dagger H V_i|\psi\rangle$. For this, it suffices to estimate each $\langle\psi|V_i^\dagger H_j V_i|\psi\rangle$ individually. If H_j acts on qubits q_1, q_2, q_3 , then we simply query C_ψ for the original state of qubits q_1, q_2, q_3 , apply $V_{i,q_1} \otimes V_{i,q_2} \otimes V_{i,q_3}$ to these three qubits, and finally compute the desired expectation against H_j .
- The verification procedure now requires exponential space, since we must keep track of exponentially many cumulative 1-qubit operators V_{iq} which comprise the global i 'th operator $V_i = V_{i,1} \otimes \cdots \otimes V_{i,2^n}$.

□

8 Conclusions and open problems

In this paper, we defined a physically motivated notion of connectivity for ground spaces of quantum local Hamiltonians, and initiated its study. Specifically, we asked: Given a local Hamiltonian H and initial and final states $|\psi\rangle$ and $|\phi\rangle$, respectively, can $|\psi\rangle$ be mapped via local unitary operations to $|\phi\rangle$ through the ground space of H ? Our main results showed that the complexity of this problem can range from QCMA-complete to NEXP-complete, depending on the specific formulation of the problem. As a result, we obtained a natural QCMA-complete problem, adding to the short list of known QCMA-complete problems. To show this QCMA-hardness result, we proved the Traversal Lemma, which allows one to analyze the path a unitary evolution must take in certain settings. We further showed that the Traversal Lemma is tight up to a polynomial factor in the length of the unitary evolution considered.

We close with the following open problems. (1) References [GKMP06] and [MNPR14] show dichotomy and trichotomy theorems, respectively, for classical reconfiguration problems involving Boolean satisfiability; can similar theorems be shown in the quantum setting? For example, are there non-trivial quantum cases of GSCON which can be solved in P or BQP? (2) Our complexity theoretic results on GSCON depended crucially on the parameters m (the number of unitaries) and l (the locality of each unitary). We have shown that polynomial m and $l = 2$ characterizes QCMA, and that exponential m and $l = 1$ characterizes PSPACE. There is, however, an interesting regime left to consider: Exponential m and $l = 2$. In this case, our proof of containment in PSPACE seems to fail as each intermediate state in the evolution appears to require exponential space to represent. However, this variant of the problem is in NEXP, and we conjecture that it is in fact NEXP-complete. (3) Regarding our Traversal Lemma, can it (or some variant thereof) be used in other settings in quantum computational complexity, such as in analyzing quantum adiabatic algorithms? (4) Finally, are there other problems related to GSCON which are also complete for quantum complexity classes such as QCMA?

Acknowledgements

We thank Joel Klassen for suggesting the connection between GSCON and quantum memories, Barbara Terhal and Roberto Oliveira for helpful discussions regarding perturbation theory gadgets, Sarvagya Upadhyay, Damian Markham, Eleni Diamanti, Attila Pereszlenyi, Amer Mouawad, Vinayak Pathak, and David Gosset for insightful discussions, and the anonymous referees whose feedback helped improve this paper. Part of this work was completed during the UC Berkeley Simons Institute for the Theory of Computing's program on Quantum Hamiltonian Complexity.

SG acknowledges support from a Government of Canada NSERC Banting Postdoctoral Fellowship and the Simons Institute for the Theory of Computing at UC Berkeley.

JS acknowledges support from a Government of Canada NSERC Postdoctoral Fellowship, the French National Research Agency (ANR-09-JCJC-0067-01), and the European Union (ERC project QCC 306537). Research at the Centre for Quantum Technologies at the National University of Singapore is partially funded by the Singapore Ministry of Education and the National Research Foundation, also through the Tier 3 Grant “Random numbers from quantum processes,” (MOE2012-T3-1-009).

References

- [Amb14] A. Ambainis. On physical problems that are slightly more difficult than QMA. In *Proceedings of 29th IEEE Conference on Computational Complexity (CCC 2014)*, pages 32–43, 2014.
- [AN02] D. Aharonov and T. Naveh. Quantum NP - A survey. Available at arXiv.org e-Print quant-ph/0210077v1, 2002.
- [BB13] M. Bonamy and N. Bousquet. Recoloring bounded treewidth graphs. In *Proceedings of the 7th Latin-American Algorithms, Graphs, and Optimization Symposium (LAGOS)*, 2013.
- [BC09] P. Bonsma and L. Cereceda. Finding paths between graph colourings: PSPACE-completeness and superpolynomial distances. *Theoretical Computer Science*, 410(50):5215–5226, 2009.
- [BCWdW01] H. Buhrman, R. Cleve, J. Watrous, and R. de Wolf. Quantum fingerprinting. *Physical Review Letters*, 87(16):167902, 2001.
- [BFS11] B. Brown, S. Flammia, and N. Schuch. Computational difficulty of computing the density of states. *Physical Review Letters*, 104:040501, 2011.
- [BJL⁺11] M. Bonamy, M. Johnson, I. Lignos, V. Patel, and D. Paulusma. On the diameter of reconfiguration graphs for vertex colourings. *Electronic Notes in Discrete Mathematics*, 38(0):161–166, 2011.
- [Bon12] P. Bonsma. The complexity of rerouting shortest paths. In *Lecture Notes in Computer Science*, volume 7464, pages 222–233, 2012.
- [BR04] P. Beame and C. Re. Lecture 12: Probabilistically checkable proofs, 2004. Available at: <http://courses.cs.washington.edu/courses/cse532/04sp/lect12.pdf>.
- [BV05] S. Bravyi and M. Vyalyi. Commutative version of the local Hamiltonian problem and common eigenspace problem. *Quantum Information & Computation*, 5(3):187–215, 2005.
- [CM13] T. Cubitt and A. Montanaro. Complexity classification of local hamiltonian problems. Available at arXiv.org e-Print quant-ph/1311.3161, 2013.
- [Coo72] S. Cook. The complexity of theorem proving procedures. In *Proceedings of the 3rd ACM Symposium on Theory of Computing (STOC 1972)*, pages 151–158, 1972.

- [CvdHJ08] L. Cereceda, J. van den Heuvel, and M. Johnson. Connectedness of the graph of vertex-colorings. *Discrete Mathematics*, 308(56):913–919, 2008.
- [CvdHJ11] L. Cereceda, J. van den Heuvel, and M. Johnson. Finding paths between 3-colorings. *Journal of Graph Theory*, 67(1):69–82, 2011.
- [FHHH11] G. Fricke, S. M. Hedetniemi, S. T. Hedetniemi, and K. R. Hutson. γ -graphs of graphs. *Discussiones Mathematicae Graph Theory*, 31(3):517–531, 2011.
- [Gha13] S. Gharibian. *Approximation, proof systems, and correlations in a quantum world*. PhD thesis, University of Waterloo, 2013. Preprint at arXiv:quant-ph/1301.2632.
- [GHLS14] S. Gharibian, Y. Huang, Z. Landau, and S. W. Shin. Quantum Hamiltonian complexity. Available at arXiv.org e-Print quant-ph/1401.3916v1, 2014.
- [GK12] S. Gharibian and J. Kempe. Hardness of approximation for quantum problems. In *Proceedings of 39th International Colloquium on Automata, Languages and Programming (ICALP 2012)*, pages 387–398, 2012.
- [GKMP06] P. Gopalan, P. Kolaitis, E. Maneva, and C. Papadimitriou. The connectivity of Boolean satisfiability: Computational and structural dichotomies. In *Proceedings of the 33rd International Colloquium on Automata, Languages, and Programming (ICALP 2006)*, pages 346–357, 2006.
- [GLSW14] S. Gharibian, Z. Landau, S. W. Shin, and G. Wang. Tensor network non-zero testing. Available at arXiv.org e-Print quant-ph/1406.5279, 2014.
- [Got97] D. Gottesman. Stabilizer codes and quantum error correction. Available at arXiv.org e-Print quant-ph/9705052, 1997.
- [HJ90] R. A. Horn and C. H. Johnson. *Matrix Analysis*. Cambridge University Press, 1990.
- [IDH⁺11] T. Ito, E. D. Demaine, N. J. A. Harvey, C. H. Papadimitriou, M. Sideri, R. Uehara, and Y. Uno. On the complexity of reconfiguration problems. *Theoretical Computer Science*, 412(12–14):1054–1065, 2011.
- [IKD12] T. Ito, M. Kamiński, and E. D. Demaine. Reconfiguration of list edge-colorings in a graph. *Discrete Applied Mathematics*, 160(15):2199–2207, 2012.
- [IKOZ12] T. Ito, K. Kawamura, H. Ono, and X. Zhou. Reconfiguration of list $L(2,1)$ -labelings in a graph. In *Proceedings of the 23rd International Symposium on Algorithms and Computation*, pages 34–43, 2012.
- [JW06] D. Janzing and P. Wocjan. BQP-complete problems concerning mixing properties of classical random walks on sparse graphs. Available at arXiv.org e-Print quant-ph/0610235v2, 2006.
- [Kit01] A. Kitaev. Unpaired majorana fermions in quantum wires. *Physics-Uspekhi*, 44:131, 2001.
- [Kit03] A. Kitaev. Fault-tolerant quantum computation by anyons. *Annals of Physics*, 303(1):2–30, 2003.

- [KL09] A. Kitaev and C. Laumann. Topological phases and quantum computation. Available at arXiv.org e-Print quant-ph/0904.2771, 2009.
- [KMM12] M. Kamiński, P. Medvedev, and M. Milanic. Complexity of independent set reconfigurability problems. *Theoretical Computer Science*, 439:9–15, 2012.
- [KR03] J. Kempe and O. Regev. 3-local Hamiltonian is QMA-complete. *Quantum Information & Computation*, 3(3):258–264, 2003.
- [KSV02] A. Kitaev, A. Shen, and M. Vyalii. *Classical and Quantum Computation*. American Mathematical Society, 2002.
- [Lev73] L. Levin. Universal search problems. *Problems of Information Transmission*, 9(3):265–266, 1973.
- [MNPR14] A. Mouawad, N. Nishimura, V. Pathak, and V. Raman. Shortest reconfiguration paths in the solution space of Boolean formulas. Available at arXiv.org e-Print cs.CC/1404.3801v2, 2014.
- [MNR⁺13] A. E. Mouawad, N. Nishimura, V. Raman, N. Simjour, and A. Suzuki. On the parameterized complexity of reconfiguration problems. In *Proceedings of the 8th International Symposium on Parameterized and Exact Computation (IPEC)*, pages 281–294, 2013.
- [MNR14] A. Mouawad, N. Nishimura, and V. Raman. Vertex cover reconfiguration and beyond. Available at arXiv.org e-Print cs.CC/1402.4926, 2014.
- [NC00] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.
- [Osb12] T. J. Osborne. Hamiltonian complexity. *Reports on Progress in Physics*, 75(2):022001, 2012.
- [OT08] R. Oliveira and B. M. Terhal. The complexity of quantum spin systems on a two-dimensional square lattice. *Quantum Information & Computation*, 8(10):0900–0924, 2008.
- [PGA⁺11] M. Piani, S. Gharibian, G. Adesso, J. Calsamiglia, P. Horodecki, and A. Winter. All non-classical correlations can be activated into distillable entanglement. *Physical Review Letters*, 106:220403, 2011.
- [Sav70] W. J. Savitch. Relationships between nondeterministic and deterministic tape complexities. *Journal of Computer and System Sciences*, 4(2):177–192, 1970.
- [Sch78] T. J. Schaefer. The complexity of satisfiability problems. In *Proceedings of the 10th Symposium on Theory of computing*, pages 216–226, 1978.
- [Sch13] K. W. Schwerdtfeger. A computational trichotomy for connectivity of boolean satisfiability. *CoRR*, 2013. Available at arXiv.org e-Print abs/1312.4524.
- [Wat08] J. Watrous. Lecture 2: Mathematical Preliminaries Part II, 2008. Latest version available at: www.cs.uwaterloo.ca/~watrous/CS766/.

- [Wat09] J. Watrous. *Encyclopedia of Complexity and System Science*, chapter Quantum Computational Complexity. Springer, 2009.
- [Wil13] M. W. Wilde. *Quantum Information Theory*. Cambridge University Press, 2013.
- [Win99] A. Winter. Coding theorem and strong converse for quantum channels. 45(7):2481–2485, 1999.
- [WJB03] P. Wocjan, D. Janzing, and T. Beth. Two QCMA-complete problems. *Quantum Information & Computation*, 3(6):635–643, 2003.
- [WY08] P. Wocjan and J. Yard. The Jones polynomial: quantum algorithms and applications in quantum complexity theory. *Quantum Information & Computation*, 8(1):147–180, 2008.

A Proofs for Section 3

Proof of Lemma 3.1. Any 2×2 unitary U can be written in terms of parameters $0 \leq x \leq 1$ and $0 \leq \phi_1, \phi_2, \phi_3 \leq 2\pi$ such that

$$U = \begin{pmatrix} \sqrt{x} e^{i\phi_1} & \sqrt{1-x} e^{i\phi_2} \\ \sqrt{1-x} e^{i\phi_3} & \sqrt{x} e^{i\phi_4} \end{pmatrix}, \quad (17)$$

where we let $\phi_4 := -\phi_1 + \phi_2 + \phi_3 + \pi$ for brevity. The net is constructed by a straightforward discretization of the ranges of x , ϕ_1 , ϕ_2 , and ϕ_3 into segments of size $\delta > 0$, for δ to be chosen as needed. For any unitary U , there hence exist parameters $0 \leq y \leq 1$ and $0 \leq \theta_1, \theta_2, \theta_3 \leq 2\pi$ in the discretization such that $|x - y|, |\phi_1 - \theta_1|, |\phi_2 - \theta_2|, |\phi_3 - \theta_3| \leq \delta$. We now upper bound $\|U - \tilde{U}\|_\infty$, where we have defined the unitary matrix

$$\tilde{U} := \begin{pmatrix} \sqrt{y} e^{i\theta_1} & \sqrt{1-y} e^{i\theta_2} \\ \sqrt{1-y} e^{i\theta_3} & \sqrt{y} e^{i\theta_4} \end{pmatrix}$$

with $\theta_4 := -\theta_1 + \theta_2 + \theta_3 + \pi$ (implying $|\phi_4 - \theta_4| \leq 3\delta$). We first upper bound the magnitude of each entry of $U - \tilde{U}$ individually. For $j \in \{1, 4\}$, we have

$$\begin{aligned} |\sqrt{x} e^{i\phi_j} - \sqrt{y} e^{i\theta_j}| &\leq |\sqrt{x} e^{i\phi_j} - \sqrt{x} e^{i\theta_j}| + |\sqrt{x} e^{i\theta_j} - \sqrt{y} e^{i\theta_j}| \\ &= \sqrt{x} |e^{i\phi_j} - e^{i\theta_j}| + |\sqrt{x} - \sqrt{y}| \\ &\leq |\phi_j - \theta_j| + \sqrt{|x - y|} \\ &\leq 4\sqrt{\delta}, \end{aligned}$$

where we used the facts that $x \leq 1$, $|e^{i\phi} - e^{i\theta}| \leq |\phi - \theta|$ when $|\phi - \theta| \leq 1$, and the inequality $|\sqrt{x} - \sqrt{y}| \leq \sqrt{|x - y|}$. The same argument yields $|\sqrt{1-x} e^{i\phi_j} - \sqrt{1-y} e^{i\theta_j}| \leq 4\sqrt{\delta}$ for $j \in \{2, 3\}$.

We now use our bounds on each entry of $U - \tilde{U}$ as follows. For $\|A\|_{\max} := \max_{ij} |A_{ij}|$, it holds that $\|A\|_\infty \leq d \|A\|_{\max}$ for any $A \in L(\mathbb{C}^d)$ (see, e.g., [HJ90]). Hence,

$$\|U - \tilde{U}\|_\infty \leq 8\sqrt{\delta}.$$

Thus, in order to obtain an ϵ -net over single-qubit unitaries, it suffices to set $\delta = \epsilon^2/64$.

To complete the proof of our claim, we now need to bound the size of our net. Since we have 4 parameters $x, \phi_1, \phi_2, \phi_3$, each discretized into segments of length $\delta \in O(\epsilon^2)$, our net contains $O(\epsilon^{-8})$ elements. Ordering our net elements by canonically ordering the discretization of each individual parameter $x, \phi_1, \phi_2, \phi_3$ thus implies we can represent each U_i in our net using $O(\log(1/\epsilon))$ bits and retrieve U_i in time $O(\log^2(1/\epsilon))$. \square

Proof of Lemma 3.3. The construction of N is straightforward: Cast a δ -net over the unit disk for each entry (i, j) of a $d \times d$ complex matrix, for δ to be chosen as needed. For the checking algorithm C , let $|u_i\rangle$ denote the i 'th column of $\tilde{U} \in N$. Then, defining $B := \sum_{i=1}^d |u_i\rangle\langle u_i|$, C accepts if and only if

$$\|B - I\|_\infty \leq \frac{\epsilon}{2(d + \epsilon)}. \quad (18)$$

Finally, the rounding algorithm R maps input $\tilde{U} \in N$ to a matrix U whose i 'th column is given by $|u'_i\rangle := B^{-1/2} |u_i\rangle$. We remark that the rounding algorithm is heavily inspired by the epsilon net construction in [PGA⁺11, Gha13].

In order to proceed with the proof, we require a δ' -net D' over d -dimensional vectors, where $\delta' := \epsilon/[6d(d + \epsilon)]$. For this, let D denote our δ -net cast over the unit disk in our construction of N , and set $\delta := \delta'/\sqrt{d}$. Then, we claim that $D' := D^{\times d}$ gives us the desired δ' -net over \mathbb{C}^d . To see this, for $\mathbf{v} \in \mathbb{C}^d$, let \mathbf{w} be the vector obtained by snapping the coordinates of \mathbf{v} to the δ -net. Then,

$$\|\mathbf{v} - \mathbf{w}\|_2 = \sqrt{\sum_{i=1}^d (v_i - w_i)^2} \leq \sqrt{d\delta^2} = \delta'.$$

We now prove that N is an ϵ -pseudo-net. Let $U \in \mathcal{U}(\mathbb{C}^d)$. We first show that there exists $\tilde{U} \in N$ such that C accepts \tilde{U} , and that $\|U - \tilde{U}\|_\infty \leq \epsilon$. We proceed as follows: For each column $|u_i\rangle$ of U , replace it with a δ' -close vector $|u'_i\rangle \in D'$. Letting \tilde{U} denote the resulting matrix, note that $\tilde{U} \in N$. We now show the required two properties:

1. (\tilde{U} is accepted by C) Let $A := \sum_{i=1}^d |u'_i\rangle\langle u'_i|$. Then,

$$\|A - I\|_\infty \leq \sum_{i=1}^d \||u'_i\rangle\langle u'_i| - |u_i\rangle\langle u_i|\|_\infty \leq \sum_{i=1}^d \||u'_i\rangle\langle u'_i| - |u_i\rangle\langle u_i|\|_F \leq (2 + \delta')d\delta' \leq 3d\delta', \quad (19)$$

where the last inequality follows since $\delta' \leq 1$, and the third inequality follows from Equation (3) and the fact that $\||u'_i\rangle\|_2 \leq \||u_i\rangle\|_2 + \||u'_i\rangle - |u_i\rangle\|_2 \leq \delta' + 1$. Thus, \tilde{U} is accepted by C since

$$\|A - I\|_\infty \leq 3d\delta' = \frac{\epsilon}{2(d + \epsilon)}.$$

2. ($\|U - \tilde{U}\|_\infty \leq \epsilon$) We have

$$\|U - \tilde{U}\|_\infty \leq \sum_{i=1}^d \||u_i\rangle\langle i| - |u'_i\rangle\langle i|\|_\infty = \sum_{i=1}^d \||u_i\rangle - |u'_i\rangle\|_2 \leq d\delta' \leq \epsilon, \quad (20)$$

where the second inequality holds since D' is a δ' -net.

Conversely, suppose that $\tilde{U} \in N$. We show that if \tilde{U} is accepted by C , then R maps \tilde{U} to a unitary $U \in \mathcal{U}(\mathbb{C}^d)$ such that $\|\tilde{U} - U\|_\infty \leq \epsilon$. To do this, we first show that B (as used in Equation (18)) is invertible (otherwise, the algorithm R we have described is not well-defined). Indeed, suppose to the contrary that $B|v\rangle = 0$ for unit vector $|v\rangle$. Then, $\|(B - I)|v\rangle\|_2 = 1$. But this contradicts the fact that C accepts \tilde{U} , i.e., $\|B - I\|_\infty \leq \epsilon/[2(d + \epsilon)] < 1$. Next, observe that U is unitary since

$$\sum_{i=1}^d |u'_i\rangle\langle u'_i| = \sum_{i=1}^d B^{-1/2} |u_i\rangle\langle u_i| B^{-1/2} = B^{-1/2} B B^{-1/2} = I.$$

Finally, to show that $\|\tilde{U} - U\|_\infty \leq \epsilon$, by the same argument as in Equation (20), we have

$$\|U - \tilde{U}\|_\infty \leq \sum_{i=1}^d \| |u_i\rangle - |u'_i\rangle \|_2 = \sum_{i=1}^d \left\| (I - B^{-1/2}) |u_i\rangle \right\|_2 \leq d \|I - B^{-1/2}\|_\infty. \quad (21)$$

Thus, we are left to upper bound $\|I - B^{-1/2}\|_\infty$. We instead first upper bound $\|I - B\|_\infty$; using an argument analogous to Equation (19), we have that $\|I - B\|_\infty \leq 3d\delta'$. Applying now the fact that if $x \neq 0$ and $|x - 1| \leq y$, then $|(1/\sqrt{x}) - 1| \leq y/(1 - y)$, it follows that $\|I - B^{-1/2}\|_\infty \leq (3d\delta')/(1 - 3d\delta')$. Substituting this bound into Equation (21), we conclude that $\|U - \tilde{U}\|_\infty \leq \epsilon$. This completes the proof that N constitutes an ϵ -pseudo-net.

Next, to bound the size of the net N , note that since $\delta \in \Theta(\epsilon/d^{5/2})$, a trivial construction of a δ -net over the unit disk (i.e., place a square lattice of points down on the unit disk) has $O(d^5/\epsilon^2)$ elements. Since we cast the δ -net over d^2 matrix entries, the size of N is $O(d^7/\epsilon^2)$.

Finally, to compute \tilde{U}_i given i using $O(d^2 \log^2(d^{5/2}/\epsilon))$ bit operations, note that i encodes the entries of d^2 matrix positions (s, t) of \tilde{U}_i , each of which requires $\log(d^{5/2}/\epsilon)$ bits³ to encode which element from the δ -net we have at position (s, t) . Since \tilde{U}_i has d^2 entries which need to be computed given i , the claim follows. \square

³Simply encode the offsets on the imaginary and real axes.